

A Generic Logic for Proving Linearizability (Extended Version)

Artem Khyzha¹, Alexey Gotsman¹, and Matthew Parkinson²

¹ IMDEA Software Institute

² Microsoft Research Cambridge

Abstract. Linearizability is a commonly accepted notion of correctness for libraries of concurrent algorithms, and recent years have seen a number of proposals of program logics for proving it. Although these logics differ in technical details, they embody similar reasoning principles. To explicate these principles, we propose a logic for proving linearizability that is generic: it can be instantiated with different means of compositional reasoning about concurrency, such as separation logic or rely-guarantee. To this end, we generalise the Views framework for reasoning about concurrency to handle relations between programs, required for proving linearizability. We present sample instantiations of our generic logic and show that it is powerful enough to handle concurrent algorithms with challenging features, such as helping.

1 Introduction

To manage the complexity of constructing concurrent software, programmers package often-used functionality into *libraries* of concurrent algorithms. These encapsulate data structures, such as queues and lists, and provide clients with a set of methods that can be called concurrently to operate on these (e.g., `java.util.concurrent`). To maximise performance, concurrent libraries may use sophisticated non-blocking techniques, allowing multiple threads to operate on the data structure with minimum synchronisation. Despite this, each library method is usually expected to behave as though it executes atomically. This requirement is formalised by the standard notion of correctness for concurrent libraries, *linearizability* [14], which establishes a form of a simulation between the original *concrete* library and another *abstract* library, where each method is implemented atomically.

A common approach to proving linearizability is to find a *linearization point* for every method of the concrete library at which it can be thought of taking effect.³ Given an execution of a concrete library, the matching execution of the abstract library, required to show the simulation, is constructed by executing the atomic abstract method at the linearization point of the concrete method. A difficulty in this approach is that linearization points are often not determined by a statically chosen point in the method code. For example, in concurrent algorithms with *helping* [13], a method may execute an operation originally requested by another method, called in a different thread; then the linearization point of the latter method is determined by an action of the former.

³ Some algorithms cannot be reasoned about using linearization points, which we discuss in §7.

Recent years have seen a number of program logics for proving linearizability (see [6] for a survey). To avoid reasoning about the high number of possible interleavings between concurrently executing threads, these logics often use *thread-modular* reasoning. They establish protocols that threads should follow when operating on the shared data structure and reason separately about every thread, assuming that the rest follow the protocols. The logics for proving linearizability, such as [26,18], usually borrow thread-modular reasoning rules from logics originally designed for proving non-relational properties of concurrent programs, such as rely-guarantee [15], separation logic [21] or combinations thereof [26,7]. Although this leads the logics to differ in technical details, they use similar methods for reasoning about linearizability, usually based on linearization points. Despite this similarity, designing a logic for proving linearizability that uses a particular thread-modular reasoning method currently requires finding the proof rules and proving their soundness afresh.

To consolidate this design space of linearization-point-based reasoning, we propose a logic for linearizability that is *generic*, i.e., can be instantiated with different means of thread-modular reasoning about concurrency, such as separation logic [21] or rely-guarantee [15]. To this end, we build on the recently-proposed Views framework [3], which unifies thread-modular logics for concurrency, such as the above-mentioned ones. Our contribution is to generalise the framework to reason about relations between programs, required for proving linearizability. In more detail, assertions in our logic are interpreted over a monoid of *relational views*, which describe relationships between the states of the concrete and the abstract libraries and the protocol that threads should follow in operating on these. The operation of the monoid, similar to the separating conjunction in separation logic [21], combines the assertions in different threads while ensuring that they agree on the protocols of access to the state. The choice of a particular relational view monoid thus determines the thread-modular reasoning method used by our logic.

To reason about linearization points, relational views additionally describe a set of special *tokens* (as in [26,18,2]), each denoting a one-time permission to execute a given atomic command on the state of the abstract library. The place where this permission is used in the proof of a concrete library method determines its linearization point, with the abstract command recorded by the token giving its specification. Crucially, reasoning about the tokens is subject to the protocols established by the underlying thread-modular reasoning method; in particular, their *ownership* can be transferred between different threads, which allows us to deal with helping.

We prove the soundness of our generic logic under certain conditions on its instantiations (Definition 2, §3). These conditions represent our key technical contribution, as they capture the essential requirements for soundly combining a given thread-modular method for reasoning about concurrency with the linearization-point method for reasoning about linearizability.

To illustrate the use of our logic, we present its example instantiations where thread-modular reasoning is done using disjoint concurrent separation logic [20] and a combination of separation logic and rely-guarantee [26]. We then apply the latter instantiation to prove the correctness of a sample concurrent algorithm with helping. We expect that our results will make it possible to systematically design logics using the plethora of

$$\begin{array}{c}
\longrightarrow \subseteq \text{Com} \times \text{PCom} \times \text{Com} \\
\frac{C_1 \xrightarrow{\alpha} C'_1}{C_1 ; C_2 \xrightarrow{\alpha} C'_1 ; C_2} \quad \frac{}{C^* \xrightarrow{\text{id}} C ; C^*} \quad \frac{}{\alpha \xrightarrow{\alpha} \text{skip}} \\
\frac{i \in \{1,2\}}{C_1 + C_2 \xrightarrow{\text{id}} C_i} \quad \frac{}{\text{skip} ; C \xrightarrow{\text{id}} C} \quad \frac{}{C^* \xrightarrow{\text{id}} \text{skip}} \\
\longrightarrow \subseteq (\text{Com} \times \text{State}) \times (\text{ThreadID} \times \text{PCom}) \times (\text{Com} \times \text{State}) \\
\frac{\sigma' \in \llbracket \alpha \rrbracket_t(\sigma) \quad C \xrightarrow{\alpha} C'}{\langle C, \sigma \rangle \xrightarrow{t, \alpha} \langle C', \sigma' \rangle}
\end{array}$$

Fig. 1. The operational semantics of sequential commands

other methods for thread-modular reasoning that have been shown to be expressible in the Views framework [20,4,1].

2 Methods Syntax and Sequential Semantics

We consider concurrent programs that consist of two components, which we call *libraries* and *clients*. Libraries provide clients with a set of methods, and clients call them concurrently. We distinguish *concrete* and *abstract* libraries, as the latter serve as specification for the former due to its methods being executed atomically.

Syntax. Concrete methods are implemented as *sequential commands* having the syntax:

$$C \in \text{Com} ::= \alpha \mid C ; C \mid C + C \mid C^* \mid \text{skip}, \quad \text{where } \alpha \in \text{PCom}$$

The grammar includes *primitive commands* α from a set PCom, sequential composition $C ; C$, non-deterministic choice $C + C$ and a finite iteration C^* (we are interested only in terminating executions) and a termination marker skip. We use $+$ and $(\cdot)^*$ instead of conditionals and while loops for theoretical simplicity: as we show at the end of this section, given appropriate primitive commands the conditionals and loops can be encoded. We also assume a set ACom of *abstract primitive commands*, ranged over by A, with which we represent methods of an abstract library.

Semantics. We assume a set State of concrete states of the memory, ranged over by σ , and abstract states AState, ranged over by Σ . The memory is shared among N threads with thread identifiers ThreadID = $\{1, 2, \dots, N\}$, ranged over by t .

We assume that semantics of each primitive command α is given by a non-deterministic *state transformer* $\llbracket \alpha \rrbracket_t : \text{State} \rightarrow \mathcal{P}(\text{State})$, where $t \in \text{ThreadID}$. For a state σ , the set of states $\llbracket \alpha \rrbracket_t(\sigma)$ is the set of possible resulting states for α executed atomically in a state σ and a thread t . State transformers may have different semantics depending on a thread identifier, which we use to introduce thread-local memory cells later in the technical development. Analogously, we assume semantics of abstract primitive commands with state transformers $\llbracket A \rrbracket_t : \text{AState} \rightarrow \mathcal{P}(\text{AState})$, all of which update abstract states atomically. We also assume a primitive command $\text{id} \in \text{PCom}$ with the interpretation $\llbracket \text{id} \rrbracket_t(\sigma) \triangleq \{\sigma\}$, and its abstract counterpart $\text{id} \in \text{ACom}$.

The sets of primitive commands PCom and ACom as well as corresponding state transformers are parameters of our framework. In Figure 1 we give rules of operational

semantics of sequential commands, which are parametrised by semantics of primitive commands. That is, we define a transition relation $\longrightarrow \subseteq (\text{Com} \times \text{State}) \times (\text{ThreadID} \times \text{PCom}) \times (\text{Com} \times \text{State})$, so that $\langle C, \sigma \rangle \xrightarrow{t, \alpha} \langle C', \sigma' \rangle$ indicates a transition from C to C' updating the state from σ to σ' with a primitive command α in a thread t . The rules of the operational semantics are standard.

Let us show how to define traditional control flow primitives, such as an if-statement and a while-loop, in our programming language. Assuming a language for arithmetic expressions, ranged over by E , and a function $\llbracket E \rrbracket_\sigma$ that evaluates expressions in a given state σ , we define a primitive command $\text{assume}(E)$ that acts as a filter on states, choosing only those where E evaluates to non-zero values.

$$\llbracket \text{assume}(E) \rrbracket_t(\sigma) \triangleq \text{if } \llbracket E \rrbracket_\sigma \neq 0 \text{ then } \{\sigma\} \text{ else } \emptyset.$$

Using $\text{assume}(E)$ and the C-style negation $!E$ in expressions, a conditional and a while-loop can be implemented as the following commands:

$$\begin{aligned} \text{if } E \text{ then } C_1 \text{ else } C_2 &\triangleq (\text{assume}(E); C_1) + (\text{assume}(!E); C_2) \\ \text{while } E \text{ do } C &\triangleq (\text{assume}(E); C)^*; \text{assume}(!E) \end{aligned}$$

3 The generic logic

In this section, we present our framework for designing program logics for linearizability proofs. Given a concrete method and a corresponding abstract method, we aim to demonstrate that the former has a linearization point either within its code or in the code of another thread. The idea behind such proofs is to establish simulation between concrete and abstract methods using linearization points to determine when the abstract method has to make a transition to match a given execution of the concrete method. To facilitate such simulation-based proofs, we design our relational logic so that formulas in it denote relations between concrete states, abstract states and special *tokens*.

Tokens are our tool for reasoning about linearization points. At the beginning of its execution in a thread t , each concrete method m is given a token $\text{todo}(A_m)$ of the corresponding abstract primitive command A_m . The token represents a one-time permission for the method to take effect, i.e. to perform a primitive command A_m on an abstract machine. When the permission is used, a token $\text{todo}(A_m)$ in a thread t is irreversibly replaced with $\text{done}(A_m)$. Thus, by requiring that a method start its execution in a thread t with a token $\text{todo}(A_m)$ and ends with $\text{done}(A_m)$, we ensure that it has in its code a linearization point. The tokens of all threads are described by $\Delta \in \text{Tokens}$:

$$\text{Tokens} = \text{ThreadID} \rightarrow (\{\text{todo}(A) \mid A \in \text{APCom}\} \cup \{\text{done}(A) \mid A \in \text{APCom}\})$$

Reasoning about states and tokens in the framework is done with the help of relational *views*. We assume a set Views , ranged over by p, q and r , as well as a reification function $\llbracket _ \rrbracket : \text{Views} \rightarrow \mathcal{P}(\text{State} \times \text{AState} \times \text{Tokens})$ that interprets views as ternary relations on concrete states, abstract states and indexed sets of tokens.

Definition 1 *A relational view monoid is a commutative monoid $(\text{Views}, *, u)$, where Views is an underlying set of relational views, $*$ is a monoid operation and u is a unit.*

The monoid structure of relational views allows treating them as restrictions on the environment of threads. Intuitively, each thread uses views to declare a protocol that other threads should follow while operating with concrete states, abstract states and tokens. Similarly to the separating conjunction from separation logic, the monoid operation $*$ (view composition) applied to a pair of views combines protocols of access to the state and ensures that they do not contradict each other.

Disjoint Concurrent Separation logic. To give an example of a view monoid, we demonstrate the structure inspired by Disjoint Concurrent Separation logic (DCSL). A distinctive feature of DCSL is that its assertions enforce a protocol, according to which threads operate on disjoint pieces of memory. We assume a set of values Val , of which a subset $\text{Loc} \subseteq \text{Val}$ represents heap addresses. By letting $\text{State} = \text{AState} = (\text{Loc} \rightarrow_{\text{fin}} \text{Val}) \cup \{\bot\}$ we represent a state as either a finite partial function from locations to values or an exceptional faulting state \bot , which denotes the result of an invalid memory access. We define an operation \bullet on states, which results in \bot if either of the operands is \bot , or the union of partial functions if their domains are disjoint. Finally, we assume that the set PCom consists of standard heap-manipulating commands with usual semantics [21,3].

We consider the view monoid $(\mathcal{P}((\text{State} \setminus \{\bot\}) \times (\text{AState} \setminus \{\bot\}) \times \text{Tokens}), *_{\text{SL}}, ([\], [\], [\]))$: the unit is a triple of nowhere defined functions $[\]$, and the view composition defined as follows:

$$p *_{\text{SL}} p' \triangleq \{(\sigma \bullet \sigma', \Sigma \bullet \Sigma', \Delta \uplus \Delta') \mid (\sigma, \Sigma, \Delta) \in p \wedge (\sigma', \Sigma', \Delta') \in p'\}.$$

In this monoid, the composition enforces a protocol of exclusive ownership of parts of the heap: a pair of views can be composed only if they do not simultaneously describe the content of the same heap cell or a token. Since tokens are exclusively owned in DCSL, they cannot be accessed by other threads, which makes it impossible to express a helping mechanism with the DCSL views. In §5, we present another instance of our framework and reason about helping in it.

Reasoning about linearization points. We now introduce *action judgements*, which formalise linearization-points-based approach to proving linearizability within our framework.

Let us assume that α is executed in a concrete state σ with an abstract state Σ and a set of tokens Δ satisfying a precondition p . According to the action judgement $\alpha \Vdash_t \{p\}\{q\}$, for every update $\sigma' \in \llbracket \alpha \rrbracket_t(\sigma)$ of the concrete state, the abstract state may be changed to $\Sigma' \in \llbracket A \rrbracket_{t'}(\Sigma)$ in order to satisfy the postcondition q , provided that there is a token $\text{todo}(A)$ in a thread t' . When the abstract state Σ is changed and the token $\text{todo}(A)$ of a thread t' is used, the concrete state update corresponds to a linearization point, or to a regular transition otherwise.

$$\begin{array}{ccc} \sigma & \xrightarrow{[p * r]} & \Sigma \\ \downarrow \llbracket \alpha \rrbracket & & \downarrow \llbracket A \rrbracket \\ \sigma' & \xrightarrow{[q * r]} & \Sigma' \end{array}$$

Definition 2 *The action judgement $\alpha \Vdash_t \{p\}\{q\}$ holds, iff the following is true:*

$$\forall r, \sigma, \sigma', \Sigma, \Delta. (\sigma, \Sigma, \Delta) \in [p * r] \wedge \sigma' \in \llbracket \alpha \rrbracket_t(\sigma) \implies \exists \Sigma', \Delta'. \text{LP}^*(\Sigma, \Delta, \Sigma', \Delta') \wedge (\sigma', \Sigma', \Delta') \in [q * r],$$

where LP^* is the transitive closure of the following relation:

$$\text{LP}(\Sigma, \Delta, \Sigma', \Delta') \triangleq \exists t', A. \Sigma' \in \llbracket A \rrbracket_{t'}(\Sigma) \wedge \Delta(t') = \text{todo}(A) \wedge \Delta' = \Delta[t' : \text{done}(A)],$$

and $f[x : a]$ denotes the function such that $f[x : a](x) = a$ and for any $y \neq x$, $f[x : a](y) = f(y)$.

Note that depending on pre- and postconditions p and q , $\alpha \Vdash_t \{p\}\{q\}$ may encode a regular transition, a conditional or a standard linearization point. It is easy to see that the latter is the case only when in all sets of tokens Δ from $\llbracket p \rrbracket$ some thread t' has a `todo`-token, and in all Δ' from $\llbracket q \rrbracket$ it has a `done`-token. Additionally, the action judgement may represent a conditional linearization point of another thread, as the LP relation allows using tokens of other threads.

Action judgements have a closure property that is important for thread-modular reasoning: when $\alpha \Vdash_t \{p\}\{q\}$ holds, so does $\alpha \Vdash_t \{p * r\}\{q * r\}$ for every view r . That is, execution of α and a corresponding linearization point preserves every view r that p can be composed with. Consequently, when in every thread action judgements hold of primitive commands and thread's views, all threads together mutually agree on each other's protocols of the access to the shared memory encoded in their views. This enables reasoning about every thread in isolation with the assumption that its environment follows its protocol. Thus, the action judgements formalise the requirements that instances of our framework need to satisfy in order to be sound. In this regard action judgements are inspired by semantic judgements of the Views Framework [3]. Our technical contribution is in formulating the essential requirements for thread-modular reasoning about linearizability of concurrent libraries with the linearization-point method and in extending the semantic judgement with them.

We let a *repartitioning implication* of views p and q , written $p \Rightarrow q$, denote $\forall r. \llbracket p * r \rrbracket \subseteq \llbracket q * r \rrbracket$. A repartitioning implication $p \Rightarrow q$ ensures that states satisfying p also satisfy q and additionally requires this property to preserve any view r .

Program logic. We are now in a position to present our generic logic for linearizability proofs via the linearization-point method. Assuming a view monoid and reification function as parameters, we define a minimal language Assn for assertions \mathcal{P} and \mathcal{Q} denoting sets of views:

$$\mathcal{P}, \mathcal{Q} \in \text{Assn} ::= \rho \mid \mathcal{P} * \mathcal{Q} \mid \mathcal{P} \vee \mathcal{Q} \mid \mathcal{P} \Rightarrow \mathcal{Q} \mid \exists X. \mathcal{P} \mid \dots$$

The grammar includes view assertions ρ , a syntax VAssn of which is a parameter of the framework. Formulas of Assn may contain the standard connectives from separation logic, the repartitioning implication and the existential quantification over logical variables X , ranging over a set LVar .

Let us assume an interpretation of logical variables $\mathbf{i} \in \text{Int} = \text{LVar} \rightarrow \text{Val}$ that maps logical variables from LVar to values from a finite set Val . In Figure 2, we define a function $\llbracket \cdot \rrbracket : \text{Assn} \times \text{Int} \rightarrow \text{Views}$ that we use to interpret assertions. Interpretation of assertions is parametrised by $\llbracket \cdot \rrbracket : \text{VAssn} \times \text{Int} \rightarrow \text{Views}$. In order to interpret disjunction, we introduce a corresponding operation on views and require the following properties from it:

$$\llbracket p \vee q \rrbracket = \llbracket p \rrbracket \cup \llbracket q \rrbracket \qquad (p \vee q) * r = (p * r) \vee (q * r) \qquad (1)$$

$$\begin{aligned}
\llbracket \mathcal{P} * \mathcal{Q} \rrbracket_i &= \llbracket \mathcal{P} \rrbracket_i * \llbracket \mathcal{Q} \rrbracket_i & \llbracket \mathcal{P} \Rightarrow \mathcal{Q} \rrbracket_i &= \llbracket \mathcal{P} \rrbracket_i \Rightarrow \llbracket \mathcal{Q} \rrbracket_i \\
\llbracket \mathcal{P} \vee \mathcal{Q} \rrbracket_i &= \llbracket \mathcal{P} \rrbracket_i \vee \llbracket \mathcal{Q} \rrbracket_i & \llbracket \exists X. \mathcal{P} \rrbracket_i &= \bigvee_{n \in \text{Val}} \llbracket \mathcal{P} \rrbracket_{i[X:n]}
\end{aligned}$$

Fig. 2. Satisfaction relation for the assertion language Assn

$$\begin{array}{ll}
\text{(PRIM)} \quad \frac{\forall i. \alpha \Vdash_t \{\llbracket \mathcal{P} \rrbracket_i\} \{\llbracket \mathcal{Q} \rrbracket_i\}}{\vdash_t \{\mathcal{P}\} \alpha \{\mathcal{Q}\}} & \text{(SEQ)} \quad \frac{\vdash_t \{\mathcal{P}\} C_1 \{\mathcal{P}'\} \quad \vdash_t \{\mathcal{P}'\} C_2 \{\mathcal{Q}\}}{\vdash_t \{\mathcal{P}\} C_1 ; C_2 \{\mathcal{Q}\}} \\
\text{(FRAME)} \quad \frac{\vdash_t \{\mathcal{P}\} C \{\mathcal{Q}\}}{\vdash_t \{\mathcal{P} * \mathcal{R}\} C \{\mathcal{Q} * \mathcal{R}\}} & \text{(DISJ)} \quad \frac{\vdash_t \{\mathcal{P}_1\} C \{\mathcal{Q}_1\} \quad \vdash_t \{\mathcal{P}_2\} C \{\mathcal{Q}_2\}}{\vdash_t \{\mathcal{P}_1 \vee \mathcal{P}_2\} C \{\mathcal{Q}_1 \vee \mathcal{Q}_2\}} \\
\text{(EX)} \quad \frac{\vdash_t \{\mathcal{P}\} C \{\mathcal{Q}\}}{\vdash_t \{\exists X. \mathcal{P}\} C \{\exists X. \mathcal{Q}\}} & \text{(CHOICE)} \quad \frac{\vdash_t \{\mathcal{P}\} C_1 \{\mathcal{Q}\} \quad \vdash_t \{\mathcal{P}\} C_2 \{\mathcal{Q}\}}{\vdash_t \{\mathcal{P}\} C_1 + C_2 \{\mathcal{Q}\}} \\
\text{(ITER)} \quad \frac{\vdash_t \{\mathcal{P}\} C \{\mathcal{P}\}}{\vdash_t \{\mathcal{P}\} C^* \{\mathcal{P}\}} & \text{(CONSEQ)} \quad \frac{\mathcal{P}' \Rightarrow \mathcal{P} \quad \vdash_t \{\mathcal{P}\} C \{\mathcal{Q}\} \quad \mathcal{Q} \Rightarrow \mathcal{Q}'}{\vdash_t \{\mathcal{P}'\} C \{\mathcal{Q}'\}}
\end{array}$$

Fig. 3. Proof rules

The judgements of the program logic take the form $\vdash_t \{\mathcal{P}\} C \{\mathcal{Q}\}$. In Figure 3, we present the proof rules, which are mostly standard. Among them, the PRIM rule is noteworthy, since it incorporates the simulation-based approach to reasoning about linearization points introduced by action judgements. The FRAME rule applies the idea of local reasoning from separation logic [21] to views. The CONSEQ enables weakening a precondition or a postcondition in a proof judgement and uses repartitioning implications to ensure the thread-modularity of the weakened proof judgement.

Semantics of proof judgements. We give semantics to judgements of the program logic by lifting the requirements of action judgements to sequential commands.

Definition 3 (Safety Judgement) We define safe_t as the greatest relation such that the following holds whenever $\text{safe}_t(p, C, q)$ does:

- if $C \neq \text{skip}$, then $\forall C', \alpha. C \xrightarrow{\alpha} C' \implies \exists p'. \alpha \Vdash_t \{p\} \{p'\} \wedge \text{safe}_t(p', C', q)$,
- if $C = \text{skip}$, then $p \Rightarrow q$.

Lemma 4 $\forall t, \mathcal{P}, C, \mathcal{Q}. \vdash_t \{\mathcal{P}\} C \{\mathcal{Q}\} \implies \forall i. \text{safe}_t(\llbracket \mathcal{P} \rrbracket_i, C, \llbracket \mathcal{Q} \rrbracket_i)$.

We can understand the safety judgement $\text{safe}_t(\llbracket \mathcal{P} \rrbracket_i, C, \llbracket \mathcal{Q} \rrbracket_i)$ as an obligation to create a sequence of views $\llbracket \mathcal{P} \rrbracket_i = p_1, p_2, \dots, p_{n+1} = \llbracket \mathcal{Q} \rrbracket_i$ for each finite trace $\alpha_1, \alpha_2, \dots, \alpha_n$ of C to justify each transition with action judgements $\alpha_1 \Vdash_t \{p_1\} \{p_2\}, \dots, \alpha_n \Vdash_t \{p_n\} \{p_{n+1}\}$. Thus, when $\text{safe}_t(\llbracket \mathcal{P} \rrbracket_i, C, \llbracket \mathcal{Q} \rrbracket_i)$ holds, it ensures that every step of the machine correctly preserves a correspondence between a concrete and abstract execution. Intuitively, the safety judgement lifts the simulation between concrete and abstract primitive commands established with action judgements to the implementation and specification of a method.

In Lemma 4, we establish that the proof judgements of the logic imply the safety judgements. As a part of the proof, we show that each of the proof rules of the logic holds of safety judgements. Due to space constraints, this and other proofs are given in the extended version of the paper [17].

4 Soundness

In this section, we formulate linearizability for libraries. We also formulate the soundness theorem, in which we state proof obligations that are necessary to conclude linearizability.

Libraries. We assume a set of method names Method , ranged over by m , and consider a concrete library $\ell : \text{Method} \rightarrow ((\text{Val} \times \text{Val}) \rightarrow \text{Com})$ that maps method names to commands from Com , which are parametrised by a pair of values from Val . For a given method name $m \in \text{dom}(\ell)$ and values $a, v \in \text{Val}$, a command $\ell(m, a, v)$ is an implementation of m , which accepts a as a method argument and either returns v or does not terminate. Such an unusual way of specifying method's arguments and return values significantly simplifies further development, since it does not require modelling a call stack.

Along with the library ℓ we consider its specification in the form of an abstract library $\mathcal{L} \in \text{Method} \rightarrow ((\text{Val} \times \text{Val}) \rightarrow \text{APCom})$ implementing a set of methods $\text{dom}(\mathcal{L})$ atomically as abstract primitive commands $\{\mathcal{L}(m, a, v) \mid m \in \text{dom}(\mathcal{L})\}$ parametrised by an argument a and a return value v . Given a method $m \in \text{Method}$, we assume that a parametrised abstract primitive command $\mathcal{L}(m)$ is intended as a specification for $\ell(m)$.

Linearizability. The linearizability assumes a complete isolation between a library and its client, with interactions limited to passing values of a given data type as parameters or return values of library methods. Consequently, we are not interested in internal steps recorded in library computations, but only in the interactions of the library with its client. We record such interactions using *histories*, which are traces including only events $\text{call } m(a)$ and $\text{ret } m(v)$ that indicate an invocation of a method m with a parameter a and returning from m with a return value v , or formally:

$$h ::= \varepsilon \mid (t, \text{call } m(a)) :: h \mid (t, \text{ret } m(v)) :: h.$$

Given a library ℓ , we generate all finite histories of ℓ by considering N threads repeatedly invoking library methods in any order and with any possible arguments. The execution of methods is described by semantics of commands from § 2.

We define a *thread pool* $\tau : \text{ThreadID} \rightarrow (\text{idle} \uplus (\text{Com} \times \text{Val}))$ to characterise progress of methods execution in each thread. The case of $\tau(t) = \text{idle}$ corresponds to no method running in a thread t . When $\tau(t) = (C, v)$, to finish some method returning v it remains to execute C .

Definition 5 We let $\mathcal{H}[\ell, \sigma] = \bigcup_{n \geq 0} \mathcal{H}_n[\ell, (\lambda t. \text{idle}), \sigma]$ denote the set of all possible histories of a library ℓ that start from a state σ , where for a given thread pool τ , $\mathcal{H}_n[\ell, \tau, \sigma]$ is defined as a set of histories such that $\mathcal{H}_0[\ell, \tau, \sigma] \triangleq \{\varepsilon\}$ and:

$$\begin{aligned} \mathcal{H}_n[\ell, \tau, \sigma] \triangleq & \{((t, \text{call } m(a)) :: h) \mid a \in \text{Val} \wedge m \in \text{dom}(\ell) \wedge \tau(t) = \text{idle} \wedge \\ & \exists v. h \in \mathcal{H}_{n-1}[\ell, \tau[t : (\ell(m, a, v), v)], \sigma]\} \\ \cup & \{h \mid \exists t, \alpha, C, C', \sigma', v. \tau(t) = (C, v) \wedge \langle C, \sigma \rangle \xrightarrow{t, \alpha} \langle C', \sigma' \rangle \wedge \\ & h \in \mathcal{H}_{n-1}[\ell, \tau[t : (C', v)], \sigma']\} \\ \cup & \{((t, \text{ret } m(v)) :: h) \mid m \in \text{dom}(\ell) \wedge \tau(t) = (\text{skip}, v) \wedge \\ & h \in \mathcal{H}_{n-1}[\ell, \tau[t : \text{idle}], \sigma]\} \end{aligned}$$

Thus, we construct the set of all finite histories inductively with all threads initially idling. At each step of generation, in any idling thread t any method $m \in \text{dom}(\ell)$ may be called with any argument a and an expected return value v , which leads to adding a command $\ell(m, a, v)$ to the thread pool of a thread t . Also, any thread t , in which $\tau(t) = (C, v)$, may do a transition $\langle C, \sigma \rangle \xrightarrow{t, \alpha} \langle C', \sigma' \rangle$ changing a command in the thread pool and the concrete state. Finally, any thread that has finished execution of a method's command ($\tau(t) = (\text{skip}, v)$) may become idle by letting $\tau(t) = \text{idle}$.

We define $\mathcal{H}_n[\mathcal{L}, \mathcal{T}, \Sigma]$ analogously and let the set of all histories of an abstract library \mathcal{L} starting from the initial state Σ be $\mathcal{H}[\mathcal{L}, \Sigma] = \bigcup_{n \geq 0} \mathcal{H}_n[\mathcal{L}, (\lambda t. \text{idle}), \Sigma]$.

Definition 6 For libraries ℓ and \mathcal{L} such that $\text{dom}(\ell) = \text{dom}(\mathcal{L})$, we say that \mathcal{L} *linearizes* ℓ in the states σ and Σ , written $(\ell, \sigma) \sqsubseteq (\mathcal{L}, \Sigma)$, if $\mathcal{H}[\ell, \sigma] \subseteq \mathcal{H}[\mathcal{L}, \Sigma]$.

That is, an abstract library \mathcal{L} linearizes ℓ in the states σ and Σ , if every history of ℓ can be reproduced by \mathcal{L} . The definition is different from the standard one [14]: we use the result obtained by Gotsman and Yang [10] stating that the plain subset inclusion on the sets of histories produced by concrete and abstract libraries is equivalent to the original definition of linearizability.

Soundness w.r.t. linearizability. We now explain proof obligations that we need to show for every method m of a concrete library ℓ to conclude its linearizability. Particularly, for every thread t , argument a , return value v , and a command $\ell(m, a, v)$ we require that there exist assertions $\mathcal{P}(t, \mathcal{L}(m, a, v))$ and $\mathcal{Q}(t, \mathcal{L}(m, a, v))$, for which the following Hoare-style specification holds:

$$\vdash_t \{ \mathcal{P}(t, \mathcal{L}(m, a, v)) \} \ell(m, a, v) \{ \mathcal{Q}(t, \mathcal{L}(m, a, v)) \} \quad (2)$$

In the specification of $\ell(m, a, v)$, $\mathcal{P}(t, \mathcal{L}(m, a, v))$ and $\mathcal{Q}(t, \mathcal{L}(m, a, v))$ are assertions parametrised by a thread t and an abstract command $\mathcal{L}(m, a, v)$. We require that in a thread t of all states satisfying $\mathcal{P}(t, \mathcal{L}(m, a, v))$ and $\mathcal{Q}(t, \mathcal{L}(m, a, v))$ there be only tokens $\text{todo}(\mathcal{L}(m, a, v))$ and $\text{done}(\mathcal{L}(m, a, v))$ respectively:

$$\begin{aligned} \forall i, t, \sigma, \Sigma, \Delta, r. \\ ((\sigma, \Sigma, \Delta) \in \llbracket \mathcal{P}(t, \mathcal{L}(m, a, v)) \rrbracket_i * r] \implies \Delta(t) = \text{todo}(\mathcal{L}(m, a, v))) \\ \wedge ((\sigma, \Sigma, \Delta) \in \llbracket \mathcal{Q}(t, \mathcal{L}(m, a, v)) \rrbracket_i * r] \implies \Delta(t) = \text{done}(\mathcal{L}(m, a, v))) \end{aligned} \quad (3)$$

Together, (2) and (3) impose a requirement that a concrete and an abstract method return the same return value v . We also require that the states satisfying the assertions only differ by a token of a thread t :

$$\begin{aligned} \forall i, t, A, A', r, \Delta. (\sigma, \Sigma, \Delta[t : \text{done}(A)]) \in \llbracket \mathcal{Q}(t, A) \rrbracket_i * r] \iff \\ (\sigma, \Sigma, \Delta[t : \text{todo}(A')]) \in \llbracket \mathcal{P}(t, A') \rrbracket_i * r]. \end{aligned} \quad (4)$$

Theorem 7 For given libraries ℓ and \mathcal{L} together with states σ and Σ , $(\ell, \sigma) \sqsubseteq (\mathcal{L}, \Sigma)$ holds, if $\text{dom}(\ell) = \text{dom}(\mathcal{L})$ and (2), (3) and (4) hold for every method m , thread t and values a and v .

5 The RGSep-based Logic

In this section, we demonstrate an instance of the generic proof system that is capable of handling algorithms with helping. This instance is based on RGSep [26], which combines rely-guarantee reasoning [15] with separation logic [21].

The main idea of the logic is to partition the state into several thread-local parts (which can only be accessed by corresponding threads) and the shared part (which can be accessed by all threads). The partitioning is defined by proofs in the logic: an assertion in the code of a thread restricts its local state and the shared state. In addition, the partitioning is dynamic, meaning that resources, such as a part of a heap or a token, can be moved from the local state of a thread into the shared state and vice versa. By transferring a token to the shared state, a thread gives to its environment a permission to change the abstract state. This allows us to reason about environment helping that thread.

The RGSep-based view monoid. Similarly to DCSL, we assume that states represent heaps, i.e. that $\text{State} = \text{AState} = \text{Loc} \multimap_{\text{fin}} \text{Val} \uplus \{\perp\}$, and we denote all states but a faulting one with $\text{State}_H = \text{AState}_H = \text{Loc} \multimap_{\text{fin}} \text{Val}$. We also assume a standard set of heap-manipulating primitive commands with usual semantics.

We define views as triples consisting of three components: a predicate P and binary relations R and G . A predicate $P \in \mathcal{P}((\text{State}_H \times \text{AState}_H \times \text{Tokens})^2)$ is a set of pairs (l, s) of local and shared parts of the state, where each part consists of concrete state, abstract state and tokens. *Guarantee* G and *rely* R are relations from $\mathcal{P}((\text{State} \times \text{AState} \times \text{Tokens})^2)$, which summarise how individual primitive commands executed by the method's thread (in case of G) and the environment (in case of R) may change the shared state. Together guarantee and rely establish a protocol that views of the method and its environment respectively must agree on each other's transitions, which allows us to reason about every thread separately without considering local state of other threads, assuming that they follow the protocol. The agreement is expressed with the help of a well-formedness condition on views of the RGSep-based monoid that their predicates must be *stable* under rely, meaning that their predicates take into account whatever changes their environment can make:

$$\text{stable}(P, R) \triangleq \forall l, s, s'. (l, s) \in P \wedge (s, s') \in R \implies (l, s') \in P.$$

A predicate that is stable under rely cannot be invalidated by any state transition from rely. Stable predicates with rely and guarantee relations form the view monoid with the underlying set of views $\text{Views}_{\text{RGSep}} = \{(P, R, G) \mid \text{stable}(P, R)\} \cup \{\perp\}$, where \perp denotes a special *inconsistent* view with the empty reification. The reification of other views simply joins shared and local parts of the state:

$$\llbracket (P, R, G) \rrbracket = \{(\sigma_l \bullet \sigma_s, \Sigma_l \bullet \Sigma_s, \Delta_l \uplus \Delta_s) \mid ((\sigma_l, \Sigma_l, \Delta_l), (\sigma_s, \Sigma_s, \Delta_s)) \in P\}.$$

Let an operation \cdot be defined on states analogously to DCSL. Given predicates P and P' , we let $P * P'$ be a predicate denoting the pairs of local and shared states in which the local state can be divided into two substates such that one of them together with the shared state satisfies P and the other together with the shared state satisfies P' :

$$P * P' \triangleq \{((\sigma_l \bullet \sigma'_l, \Sigma_l \bullet \Sigma'_l, \Delta_l \uplus \Delta'_l), s) \mid ((\sigma_l, \Sigma_l, \Delta_l), s) \in P \wedge ((\sigma'_l, \Sigma'_l, \Delta'_l), s) \in P'\}$$

$$\begin{aligned}
& \models: (\text{State} \times \text{AState} \times \text{Tokens}) \times (\text{State} \times \text{AState} \times \text{Tokens}) \times \text{Int} \times \text{Assn} \\
& ((\sigma_l, \Sigma_l, \Delta_l), (\sigma_s, \Sigma_s, \Delta_s), \mathbf{i}) \models E \mapsto F, \quad \text{iff } \sigma_l = \llbracket E \rrbracket_{\mathbf{i}} : \llbracket F \rrbracket_{\mathbf{i}}, \Sigma_l = [], \text{ and } \Delta_l = [] \\
& ((\sigma_l, \Sigma_l, \Delta_l), (\sigma_s, \Sigma_s, \Delta_s), \mathbf{i}) \models E \Rightarrow F, \quad \text{iff } \sigma_l = [], \Sigma_l = \llbracket E \rrbracket_{\mathbf{i}} : \llbracket F \rrbracket_{\mathbf{i}}, \text{ and } \Delta_l = [] \\
& ((\sigma_l, \Sigma_l, \Delta_l), (\sigma_s, \Sigma_s, \Delta_s), \mathbf{i}) \models [\text{todo}(A)]_t, \quad \text{iff } \sigma_l = [], \Sigma_l = [], \text{ and } \Delta_l = [t : \text{todo}(A)] \\
& ((\sigma_l, \Sigma_l, \Delta_l), (\sigma_s, \Sigma_s, \Delta_s), \mathbf{i}) \models [\text{done}(A)]_t, \quad \text{iff } \sigma_l = [], \Sigma_l = [], \text{ and } \Delta_l = [t : \text{done}(A)] \\
& ((\sigma_l, \Sigma_l, \Delta_l), (\sigma_s, \Sigma_s, \Delta_s), \mathbf{i}) \models \boxed{\pi}, \quad \text{iff } \sigma_l = [], \Sigma_l = [], \Delta_l = [], \text{ and} \\
& \quad \quad \quad ((\sigma_s, \Sigma_s, \Delta_s), ([], [], []), \mathbf{i}) \models \pi \\
& ((\sigma_l, \Sigma_l, \Delta_l), (\sigma_s, \Sigma_s, \Delta_s), \mathbf{i}) \models \pi * \pi', \quad \text{iff there exist } \sigma'_l, \sigma''_l, \Sigma'_l, \Sigma''_l, \Delta'_l, \Delta''_l \text{ such that} \\
& \quad \quad \quad \sigma_l = \sigma'_l \bullet \sigma''_l, \Sigma_l = \Sigma'_l \bullet \Sigma''_l, \Delta_l = \Delta'_l \uplus \Delta''_l, \\
& \quad \quad \quad ((\sigma'_l, \Sigma'_l, \Delta'_l), (\sigma_s, \Sigma_s, \Delta_s), \mathbf{i}) \models \pi, \text{ and} \\
& \quad \quad \quad ((\sigma''_l, \Sigma''_l, \Delta''_l), (\sigma_s, \Sigma_s, \Delta_s), \mathbf{i}) \models \pi'
\end{aligned}$$

Fig. 4. Satisfaction relation for a fragment of the assertion language VAssn

We now define the monoid operation $*$, which we use to compose views of different threads. When composing views (P, R, G) and (P', R', G') of the parallel threads, we require predicates of both to be immune to interference by all other threads and each other. Otherwise, the result is inconsistent:

$$(P, R, G) * (P', R', G') \triangleq \text{if } G \subseteq R' \wedge G' \subseteq R \text{ then } (P * P', R \cap R', G \cup G') \text{ else } \perp.$$

That is, we let the composition of views be consistently defined when the state transitions allowed in a guarantee of one thread are treated as environment transitions in the other thread, i.e. $G \subseteq R'$ and $G' \subseteq R$. The rely of the composition is $R \cap R'$, since the predicate $P * P'$ is guaranteed to be stable only under environment transitions described by both R and R' . The guarantee of the composition is $G \cup G'$, since other views need to take into account all state transitions either from G or from G' .

The RGSep-based program logic. We define the view assertion language VAssn that is a parameter of the proof system. Each view assertion ρ takes form of a triple $(\pi, \mathcal{R}, \mathcal{G})$, and the syntax for π is:

$$\begin{aligned}
E &::= a \mid X \mid E + E \mid \dots, \quad \text{where } X \in \text{LVar}, a \in \text{Val} \\
\pi &::= E = E \mid E \mapsto E \mid E \Rightarrow E \mid [\text{todo}(A)]_t \mid [\text{done}(A)]_t \mid \boxed{\pi} \mid \pi * \pi \mid \neg \pi \mid \dots
\end{aligned}$$

Formula π denotes a predicate of a view as defined by a satisfaction relation \models in Figure 4. There $E \mapsto E$ and $E \Rightarrow E$ denote a concrete and an abstract state describing singleton heaps. A non-boxed formula π denotes the view with the local state satisfying π and shared state unrestricted; $\boxed{\pi}$ denotes the view with the empty local state and the shared state satisfying π ; $\pi * \pi'$ the composition of predicates corresponding to π and π' . The semantics of the rest of connectives is standard. Additionally, for simplicity of presentation of the syntax, we require that boxed assertions $\boxed{\pi}$ be not nested (as opposed to preventing that in the definition).

The other components \mathcal{R} and \mathcal{G} of a view assertion are sets of *rely/guarantee actions* \mathcal{A} with the syntax: $\mathcal{A} ::= \pi \rightsquigarrow \pi'$. An action $\pi \rightsquigarrow \pi'$ denotes a change of a part of the shared state that satisfies π into one that satisfies π' , while leaving the rest of the shared state unchanged. We associate with an action $\pi \rightsquigarrow \pi'$ all state transitions from

the following set:

$$\begin{aligned} \llbracket \pi \rightsquigarrow \pi' \rrbracket = & \{ ((\sigma_s \bullet \sigma_s'', \Sigma_s \bullet \Sigma_s'', \Delta_s \uplus \Delta_s''), (\sigma'_s \bullet \sigma_s'', \Sigma'_s \bullet \Sigma_s'', \Delta'_s \uplus \Delta_s'')) \mid \\ & \exists \mathbf{i}. (([], [], []), (\sigma_s, \Sigma_s, \Delta_s), \mathbf{i}) \models \boxed{\pi} \wedge (([], [], []), (\sigma'_s, \Sigma'_s, \Delta'_s), \mathbf{i}) \models \boxed{\pi'} \} \end{aligned}$$

We give semantics to view assertions with the function $\llbracket \cdot \rrbracket$, that is defined as follows:

$$\llbracket (\pi, \mathcal{R}, \mathcal{G}) \rrbracket_{\mathbf{i}} \triangleq (\{(l, s) \mid (l, s, \mathbf{i}) \models \pi\}, \bigcup_{\mathcal{A} \in \mathcal{R}} \llbracket \mathcal{A} \rrbracket, \bigcup_{\mathcal{A} \in \mathcal{G}} \llbracket \mathcal{A} \rrbracket).$$

6 Example

In this section, we demonstrate how to reason about algorithms with helping using relational views. We choose a simple library ℓ implementing a concurrent increment and prove its linearizability with the RGSep-based logic.

The concrete library ℓ has one method `inc`, which increments the value of a shared counter `k` by the argument of the method. The specification of ℓ is given by an abstract library \mathcal{L} . The abstract command, provided by \mathcal{L} as an implementation of `inc`, operates with an abstract counter `K` as follows (assuming that `K` is initialised by zero):

```
1  $\mathcal{L}(\text{inc}, a, r)$ :    < --kabs := --kabs + a; assume(--kabs == r); >
```

That is, $\mathcal{L}(\text{inc}, a, r)$ atomically increments a counter and a command `assume(K == r)`, which terminates only if the return value r chosen at the invocation equals to the resulting value of `K`. This corresponds to how we specify methods' return values in §4.

In Figure 5, we show the pseudo-code of the implementation of a method `inc` in a C-style language along with a proof outline. The method $\ell(\text{inc}, a, r)$ takes one argument, increments a shared counter `k` by it and returns the increased value of the counter. Since `k` is shared among threads, they follow a protocol regulating the access to the counter. This protocol is based on flat combining [11], which is a synchronisation technique enabling a parallel execution of sequential operations.

The protocol is the following. When a thread t executes $\ell(\text{inc}, a, r)$, it first makes the argument of the method visible to other threads by storing it in an array `arg`, and lets `res[t] = nil` to signal to other threads its intention to execute an increment with that argument. It then spins in the loop on line 8, trying to write its thread identifier into a variable `L` with a compare-and-swap (CAS). Out of all threads spinning in the loop, the one that succeeds in writing into `L` becomes a *combiner*: it performs the increments requested by all threads with arguments stored in `arg` and writes the results into corresponding cells of the array `res`. The other threads keep spinning and periodically checking the value of their cells in `res` until a non-`nil` value appears in it, meaning that a combiner has performed the operation requested and marked it as finished. The protocol relies on the assumption that `nil` is a value that is never returned by the method. Similarly to the specification of the increment method, the implementation in Figure 5 ends with a command `assume(res[mytid()] = r)`.

The proof outline features auxiliary assertions defined in Figure 6. In the assertions we let `_` denote a value or a logical variable whose name is irrelevant. We assume that each program variable `var` has a unique location in the heap and denote it with `&var`. Values a , r and t are used in the formulas and the code as constants.

```

1 int L = 0, k = 0, arg[N], res[N]; \\ initially all res[i] ≠ nil
2
3 ℓ(inc, a, r):
4   { global * M(t) * [todo(ℒ(inc, a, r))]_t }
5   arg[mytid()] := a;
6   res[mytid()] := nil;
7   { global * true * (task_todo(t, a, r) ∨ task_done(t, a, r)) }
8   while (res[mytid()] = nil):
9     if (CAS(&L, 0, mytid())):
10      { &L ↦ t * ⊗_{j ∈ ThreadID} tinv(j) * true * (task_todo(t, a, r) ∨ task_done(t, a, r)) * kinv(⊥) }
11      for (i := 1; i ≤ N; ++i):
12        { &L ↦ t * ⊗_{j ∈ ThreadID} tinv(j) * kinv(⊥) * LI(i, t, a, r) }
13        if (res[i] = nil):
14          { ∃V, A, R. kinv(V) * LI(i, t, a, r) *
15            { &L ↦ t * ⊗_{j ∈ ThreadID} tinv(j) * true * task_todo(i, A, R) } }
16          k := k + arg[i];
17          { ∃V, A, R. &k ↦ V + A * &k ⇒ V * LI(i, t, a, r) *
18            { &L ↦ t * ⊗_{j ∈ ThreadID} tinv(j) * true * task_todo(i, A, R) } }
19          res[i] := k;
20          { ∃V, A, R. kinv(V + A) * LI(i + 1, t, a, r) *
21            { &L ↦ t * ⊗_{j ∈ ThreadID} tinv(j) * true * task_done(i, A, R) } }
22          { &L ↦ t * ⊗_{j ∈ ThreadID} tinv(j) * true * task_done(t, a, r) * kinv(⊥) }
23          L = 0;
24          assume(res[mytid()] = r);
25          { global * M(t) * [done(ℒ(inc, a, r))]_t }

```

Fig. 5. Proof outline for a flat combiner of a concurrent increment. Indentation is used for grouping commands.

We prove the following specification for $\ell(\text{inc}, a, r)$:

$$\mathcal{R}_t, \mathcal{G}_t \vdash_t \left\{ \begin{array}{c} \text{global} * M(t) * \\ [\text{todo}(\mathcal{L}(\text{inc}, a, r))]_t \end{array} \right\} \ell(\text{inc}, a, r) \left\{ \begin{array}{c} \text{global} * M(t) * \\ [\text{done}(\mathcal{L}(\text{inc}, a, r))]_t \end{array} \right\}$$

In the specification, $M(t)$ asserts the presence of $\text{arg}[t]$ and $\text{res}[t]$ in the shared state, and global is an assertion describing the shared state of all the threads. Thus, the pre- and postcondition of the specification differ only by the kind of token given to t .

$$\begin{aligned}
X \not\vdash Y &\triangleq \exists Y'. X \mapsto Y' * Y \neq Y' \\
M(t) &\triangleq \boxed{\text{true} * (\&\text{arg}[t] \mapsto _ * \&\text{res}[t] \not\vdash \text{nil})} \\
\text{task}_{\text{todo}}(t, a, r) &\triangleq \&\text{arg}[t] \mapsto a * \&\text{res}[t] \mapsto \text{nil} * [\text{todo}(\mathcal{L}(\text{inc}, a, r))]_t; \\
\text{task}_{\text{done}}(t, a, r) &\triangleq \&\text{arg}[t] \mapsto a * \&\text{res}[t] \mapsto r * r \neq \text{nil} * [\text{done}(\mathcal{L}(\text{inc}, a, r))]_t; \\
\text{kinv}(V) &\triangleq \&\text{k} \mapsto V * \&\text{K} \mapsto V \\
\text{LI}(i, t, a, r) &\triangleq \boxed{\text{true} * ((t < i \wedge \text{task}_{\text{done}}(t, a, r)) \vee (t \geq i \wedge (\text{task}_{\text{todo}}(t, a, r) \vee \text{task}_{\text{done}}(t, a, r)))} \\
\text{tinv}(i) &\triangleq \&\text{arg}[i] \mapsto _ * \&\text{res}[i] \mapsto _ \vee \text{task}_{\text{todo}}(i, _, _) \vee \text{task}_{\text{done}}(i, _, _) \\
\text{global} &\triangleq \boxed{(\&\text{L} \mapsto 0 * \text{kinv}(_) \vee \&\text{L} \not\vdash 0) * \bigotimes_{j \in \text{ThreadID}} \text{tinv}(j)},
\end{aligned}$$

Fig. 6. Auxiliary predicates. $\bigotimes_{j \in \text{ThreadID}} \text{tinv}(j)$ denotes $\text{tinv}(1) * \text{tinv}(2) * \dots * \text{tinv}(N)$

The main idea of the proof is in allowing a thread t to share the ownership of its token $[\text{todo}(\mathcal{L}(\text{inc}, a, r))]_t$ with the other threads. This enables two possibilities for t . Firstly, t may become a combiner. Then t has a linearization point on line 17 (when the loop index i equals to t). In this case t also *helps* other concurrent threads by performing their linearization points on line 17 (when $i \neq t$). The alternative possibility is that some other thread becomes a combiner and does a linearization point of t . Thus, the method has a non-fixed linearization point, as it may occur in the code of a different thread.

We further explain how the tokens are transferred. On line 6 the method performs the assignment $\text{res}[\text{mytid}()] := \text{nil}$, signalling to other threads about a task this thread is performing. At this step, the method transfers its token $[\text{todo}(\mathcal{L}(\text{inc}, a, r))]_t$ to the shared state, as represented by the assertion $\boxed{\text{true} * \text{task}_{\text{todo}}(t, a, r)}$. In order to take into consideration other threads interfering with t and possibly helping it, here and further we stabilise the assertion by adding a disjunct $\text{task}_{\text{done}}(t, a, r)$.

If a thread t gets help from other threads, then $\text{task}_{\text{done}}(t, a, r)$ holds, which implies that $\text{res}[t] \neq \text{nil}$ and t cannot enter the loop on line 8. Otherwise, if t becomes a combiner, it transfers $\text{kinv}(_)$ from the shared state to the local state of t to take over the ownership of the counters k and K and thus ensure that the access to the counter is governed by the mutual exclusion protocol. At each iteration i of the forall loop, $\text{res}[i] = \text{nil}$ implies that $\text{task}_{\text{todo}}(i, _, _)$ holds, meaning that there is a token of a thread i in the shared state. Consequently, on line 17 a thread t may use it to perform a linearization point of i .

The actions defining the guarantee relation \mathcal{G}_t of a thread t' are the following:

1. $\&\text{arg}[t] \mapsto _ * \&\text{res}[t] \not\vdash \text{nil} \rightsquigarrow \&\text{arg}[t] \mapsto a * \&\text{res}[t] \not\vdash \text{nil}$;
2. $\&\text{arg}[t] \mapsto a * \&\text{res}[t] \not\vdash \text{nil} \rightsquigarrow \text{task}_{\text{todo}}(t, a, r)$;
3. $\&\text{L} \mapsto 0 * \text{kinv}(_) \rightsquigarrow \&\text{L} \mapsto t$;
4. $\&\text{L} \mapsto t * \text{task}_{\text{todo}}(T, A, R) \rightsquigarrow \&\text{L} \mapsto t * \text{task}_{\text{done}}(T, A, R)$
5. $\&\text{L} \mapsto t \rightsquigarrow \&\text{L} \mapsto 0 * \text{kinv}(_)$
6. $\text{task}_{\text{done}}(t, a, r) \rightsquigarrow \&\text{arg}[t] \mapsto a * \&\text{res}[t] \mapsto r$

Out of them, conditions 2 and 6 specify transferring the token of a thread t to and from the shared state, and condition 4 describes using the shared token of a thread T . The rely relation of a thread t is then defined as the union of all actions from guarantee relations of other threads and an additional action for each thread $t' \in \text{ThreadID} \setminus \{t\}$

allowing the client to prepare a thread t' for a new method call by giving it a new token:
 $[\text{done}(\mathcal{L}(\text{inc}, A, R))]_{t'} \rightsquigarrow [\text{todo}(\mathcal{L}(\text{inc}, A', R'))]_{t'}.$

7 Related Work

There has been a significant amount of research on methods for proving linearizability. Due to space constraints, we do not attempt a comprehensive survey here (see [6]) and only describe the most closely related work.

The existing logics for linearizability that use linearization points differ in the thread-modular reasoning method used and, hence, in the range of concurrent algorithms that they can handle. Our goal in this paper was to propose a uniform basis for designing such logics and to formalise the method they use for reasoning about linearizability in a way independent of the particular thread-modular reasoning method used. We have only shown instantiations of our logic based on disjoint concurrent separation logic [20] and RGSep [26]. However, we expect that our logic can also be instantiated with more complex thread-modular reasoning methods, such as those based on concurrent abstract predicates [4] or islands and protocols [25].

Our notion of tokens is based on the idea of treating method specifications as resources when proving atomicity, which has appeared in various guises in several logics [26,18,2]. Our contribution is to formalise this method of handling linearization points independently from the underlying thread-modular reasoning method and to formulate the conditions for soundly combining the two (Definition 2, §3).

We have presented a logic that unifies the various logics based on linearization points with helping. However, much work still remains as this reasoning method cannot handle all algorithms. Some logics have introduced *speculative* linearization points to increase their applicability [25,18]; our approach to helping is closely related to this, and we hope could be extended to speculation. But there are still examples beyond this form of reasoning: for instance there are no proofs of the Herlihy-Wing queue [14] using linearization points (with helping and/or speculation). This algorithm can be shown linearizable using forwards/backwards simulation [14] and more recently has been shown to only require a backwards simulation [22]. But integrating this form of simulation with the more intricate notions of interference expressible in the Views framework remains an open problem.

Another approach to proving linearizability is the aspect-oriented method. This gives a series of properties of a queue [12] (or a stack [5]) implementation which imply that the implementation is linearizable. This method been applied to algorithms that cannot be handled with standard linearization-point-based methods. However, the aspect-oriented approach requires a custom theorem per data structure, which limits its applicability.

In this paper we concentrated on linearizability in its original form [14], which considers only finite computations and, hence, specifies only safety properties of the library. Linearizability has since been generalised to also specify liveness properties [9]. Another direction of future work is to generalise our logic to handle liveness, possibly building on ideas from [19].

When a library is linearizable, one can use its atomic specification instead of the actual implementation to reason about its clients [8]. Some logics achieve the same effect without using linearizability, by expressing library specifications as judgements in the logic rather than as the code of an abstract library [16,24,23]. It is an interesting direction of future work to determine a precise relationship between this method of specification and linearizability, and to propose a generic logic unifying the two.

8 Conclusion

We have presented a logic for proving the linearizability of concurrent libraries that can be instantiated with different methods for thread-modular reasoning. To this end, we have extended the Views framework [3] to reason about relations between programs. Our main technical contribution in this regard was to propose the requirement for axiom soundness (Definition 2, §3) that ensures a correct interaction between the treatment of linearization points and the underlying thread-modular reasoning. We have shown that our logic is powerful enough to handle concurrent algorithms with challenging features, such as helping. More generally, our work marks the first step towards unifying the logics for proving relational properties of concurrent programs.

References

1. R. Bornat, C. Calcagno, P. W. O’Hearn, and M. J. Parkinson. Permission accounting in separation logic. In *POPL*, 2005.
2. P. da Rocha Pinto, T. Dinsdale-Young, and P. Gardner. TaDA: A logic for time and data abstraction. In *ECOOP*, 2014.
3. T. Dinsdale-Young, L. Birkedal, P. Gardner, M. J. Parkinson, and H. Yang. Views: compositional reasoning for concurrent programs. In *POPL*, 2013.
4. T. Dinsdale-Young, M. Dodds, P. Gardner, M. J. Parkinson, and V. Vafeiadis. Concurrent abstract predicates. In *ECOOP*, 2010.
5. M. Dodds, A. Haas, and C. M. Kirsch. A scalable, correct time-stamped stack. In *POPL*, New York, NY, USA, 2015.
6. B. Dongol and J. Derrick. Verifying linearizability: A comparative survey. *arXiv CoRR*, 1410.6268, 2014.
7. X. Feng. Local rely-guarantee reasoning. In *POPL*, 2009.
8. I. Filipovic, P. W. O’Hearn, N. Rinetzky, and H. Yang. Abstraction for concurrent objects. *Theoretical Computer Science*, 2010.
9. A. Gotsman and H. Yang. Liveness-preserving atomicity abstraction. In *ICALP*, 2011.
10. A. Gotsman and H. Yang. Linearizability with ownership transfer. *LMCS*, 2013.
11. D. Hendler, I. Incze, N. Shavit, and M. Tzafrir. Flat combining and the synchronization-parallelism tradeoff. In *SPAA*, 2010.
12. T. A. Henzinger, A. Sezgin, and V. Vafeiadis. Aspect-oriented linearizability proofs. In *CONCUR*, 2013.
13. M. Herlihy and N. Shavit. *The art of multiprocessor programming*. 2008.
14. M. Herlihy and J. M. Wing. Linearizability: A correctness condition for concurrent objects. *ACM TOPLAS*, 1990.
15. C. B. Jones. Specification and design of (parallel) programs. In *IFIP Congress*, 1983.

16. R. Jung, D. Swasey, F. Sieczkowski, K. Svendsen, A. Turon, L. Birkedal, and D. Dreyer. Iris: Monoids and invariants as an orthogonal basis for concurrent reasoning. In *POPL*, 2015.
17. A. Khyzha, A. Gotsman, and M. Parkinson. A generic logic for proving linearizability (extended version). Available from <http://software.imdea.org/~artem/papers/rviews.pdf>.
18. H. Liang and X. Feng. Modular verification of linearizability with non-fixed linearization points. In *PLDI*, 2013.
19. H. Liang, X. Feng, and Z. Shao. Compositional verification of termination-preserving refinement of concurrent programs. In *LICS*, 2014.
20. P. W. O'Hearn. Resources, concurrency, and local reasoning. *Theoretical Computer Science*, 2007.
21. P. W. O'Hearn, J. C. Reynolds, and H. Yang. Local reasoning about programs that alter data structures. In *CSL*, 2001.
22. G. Schellhorn, H. Wehrheim, and J. Derrick. How to prove algorithms linearisable. In *CAV*. 2012.
23. I. Sergey, A. Nanevski, and A. Banerjee. Specifying and verifying concurrent algorithms with histories and subjectivity. In *ESOP*, 2015.
24. K. Svendsen and L. Birkedal. Impredicative concurrent abstract predicates. In *ESOP*, 2014.
25. A. J. Turon, J. Thamsborg, A. Ahmed, L. Birkedal, and D. Dreyer. Logical relations for fine-grained concurrency. In *POPL*, 2013.
26. V. Vafeiadis. Modular fine-grained concurrency verification. PhD Thesis. Technical Report UCAM-CL-TR-726, University of Cambridge, 2008.

A Additional details of the RGSep-based logic

The unit u_{RGsep} does not restrict states and the allowed state transitions of the environment, while disallowing any action in the current thread:

$$u_{\text{RGsep}} = (\{([], [], [])\} \times (\text{State} \times \text{AState} \times \text{Tokens}), (\text{State} \times \text{AState} \times \text{Tokens})^2, \emptyset)$$

Since action judgements are essential for reasoning about primitive commands in our logic, we further refine conditions under which it holds of views from the RGSep-based view monoid.

Proposition 8 *The action judgement $\alpha \Vdash_t \{(P, R, G)\}\{(Q, R, G)\}$ holds, if it is true that:*

- $\forall \sigma_l, \sigma_s, \sigma'_l, \sigma'_s, \Sigma_l, \Sigma_s, \Delta_l, \Delta_s. ((\sigma_l, \Sigma_l, \Delta_l), (\sigma_s, \Sigma_s, \Delta_s)) \in P \wedge$
 $\sigma'_l \bullet \sigma'_s \in \llbracket \alpha \rrbracket_t(\sigma_l \bullet \sigma_s) \implies \exists \Sigma'_l, \Sigma'_s, \Delta'_l, \Delta'_s. ((\sigma'_l, \Sigma'_l, \Delta'_l), (\sigma'_s, \Sigma'_s, \Delta'_s)) \in Q \wedge$
 $((\sigma_s, \Sigma_s, \Delta_s), (\sigma'_s, \Sigma'_s, \Delta'_s)) \in G \wedge \text{LP}^*(\Sigma_l \bullet \Sigma_s, \Delta_l \uplus \Delta_s, \Sigma'_l \bullet \Sigma'_s, \Delta'_l \uplus \Delta'_s);$
- $\llbracket \alpha \rrbracket_t(\sigma) \neq \perp \implies \forall \sigma'. \llbracket \alpha \rrbracket_t(\sigma \bullet \sigma') = \{\sigma'' \bullet \sigma' \mid \sigma'' \in \llbracket \alpha \rrbracket_t(\sigma)\}.$

The requirement to primitive commands in Proposition 8 is similar to that of the action judgements. The difference is that in the RG-based proof system it is not necessary to require α to preserve any view r of the environment: since a predicate P_r of any view (P_r, R_r, G_r) in another thread is stable under R_r , it is also stable under $G \subseteq R_r$ whenever $(P, R, G) * (P_r, R_r, G_r)$ is defined. Consequently, views of the environment are never invalidated by local transitions. Using the premise of Proposition 8 in PRIM rule makes it closer to the standard proof rule for the atomic step in Rely/Guarantee.

B Compositionality properties of the safety relation

In this section, we formulate and prove compositionality properties of the safety relation.

For further reference we restate the definition of a repartitioning implication:

$$p \Rightarrow q \triangleq \forall r. \lfloor p * r \rfloor \subseteq \lfloor q * r \rfloor. \quad (5)$$

Lemma 9 *The safety relation safe_t has the following closure properties:*

- FRAME: $\forall t, C, p, q, r. \text{safe}_t(p, C, q) \implies \text{safe}_t(p * r, C, q * r);$
- CHOICE: $\forall t, C_1, C_2, p, q. \text{safe}_t(p, C_1, q) \wedge \text{safe}_t(p, C_2, q) \implies \text{safe}_t(p, C_1 + C_2, q);$
- ITER: $\forall t, C, p. \text{safe}_t(p, C, p) \implies \text{safe}_t(p, C^*, p);$
- SEQ: $\forall t, C_1, C_2, p, p', q. \text{safe}_t(p, C_1, p') \wedge \text{safe}_t(p', C_2, q) \implies \text{safe}_t(p, C_1 ; C_2, q);$
- CONSEQ: $\forall t, C, p, p', q, q'. p' \Rightarrow p \wedge \text{safe}_t(p, C, q) \wedge q \Rightarrow q' \implies \text{safe}_t(p', C, q');$
- DISJ: $\forall t, C, p_1, p_2, q_1, q_2. \text{safe}_t(p_1, C, q_1) \wedge \text{safe}_t(p_2, C, q_2) \implies$
 $\text{safe}_t(p_1 \vee p_2, C, q_1 \vee q_2).$

We prove all of the properties by coinduction. To this end, we take the fixed-point definition of safe_t . We consider

$$F_t : \mathcal{P}(\text{Views} \times \text{Com} \times \text{Views}) \rightarrow \mathcal{P}(\text{Views} \times \text{Com} \times \text{Views})$$

defined as follows:

$$F_t(X) \triangleq \{(p, C, q) \mid \forall C', \alpha. C \xrightarrow{\alpha} C' \implies \exists p'. \alpha \Vdash_t \{p\}\{p'\} \wedge (p', C', q) \in X\} \\ \cup \{(p, \text{skip}, q) \mid p \Rightarrow q\}$$

Note that a powerset domain ordered by inclusion is a complete lattice and F is a mapping on it, which means that F is monotone. Consequently, by Knaster-Tarski fixed-point theorem F has the greatest fixed-point. It is easy to see that $\text{safe}_t \triangleq \text{gfp } F_t$ in Definition 3.

In the proof of Lemma 9 we use the following properties of the action judgement and the \Rightarrow relation.

Proposition 10 (Locality)

$$\forall p, q, r. p \Rightarrow q \implies p * r \Rightarrow q * r; \\ \forall t, \alpha, p, q, r. \alpha \Vdash_t \{p\}\{q\} \implies \alpha \Vdash_t \{p * r\}\{q * r\}.$$

Proposition 11 (Consequence)

$$\forall t, \alpha, p, q, p', q'. p' \Rightarrow p \wedge \alpha \Vdash_t \{p\}\{q\} \wedge q \Rightarrow q' \implies \alpha \Vdash_t \{p'\}\{q'\}.$$

The proofs of Propositions 10 and 11 are straightforward: both properties can be easily checked after unfolding definitions of action judgements.

Proposition 12 (Distributivity)

$$\forall t, \alpha, p_1, p_2, q_1, q_2. \alpha \Vdash_t \{p_1\}\{q_1\} \wedge \alpha \Vdash_t \{p_2\}\{q_2\} \implies \alpha \Vdash_t \{p_1 \vee p_2\}\{q_1 \vee q_2\}.$$

Proof. According to the Definition 2 of the action judgement $\alpha \Vdash_t \{p_1 \vee p_2\}\{q_1 \vee q_2\}$, in order to prove the latter we need to demonstrate the following:

$$\forall r, \sigma, \sigma', \Sigma, \Delta. \sigma' \in \llbracket \alpha \rrbracket_t(\sigma) \wedge (\sigma, \Sigma, \Delta) \in \lfloor (p_1 \vee p_2) * r \rfloor \implies \\ \exists \Sigma', \Delta'. \text{LP}^*(\Sigma, \Delta, \Sigma', \Delta') \wedge (\sigma', \Sigma', \Delta') \in \lfloor (q_1 \vee q_2) * r \rfloor. \quad (6)$$

Let us consider any view r , states σ, σ', Σ and tokens Δ such that both $\sigma' \in \llbracket \alpha \rrbracket_t(\sigma)$ and $(\sigma, \Sigma, \Delta) \in \lfloor (p_1 \vee p_2) * r \rfloor$ hold. According to the properties of disjunction stated in equalities (1),

$$\lfloor (p_1 \vee p_2) * r \rfloor = \lfloor (p_1 * r) \vee (p_2 * r) \rfloor = \lfloor p_1 * r \rfloor \cup \lfloor p_2 * r \rfloor.$$

Consequently, $(\sigma, \Sigma, \Delta) \in \lfloor p_1 * r \rfloor \cup \lfloor p_2 * r \rfloor$.

Let us assume that $(\sigma, \Sigma, \Delta) \in \lfloor p_1 * r \rfloor$ (the other case is analogous). Then according to the action judgement $\alpha \Vdash_t \{p_1\}\{q_1\}$, there exist Σ' and Δ' such that:

$$\text{LP}^*(\Sigma, \Delta, \Sigma', \Delta') \wedge (\sigma', \Sigma', \Delta') \in \lfloor q_1 * r \rfloor. \quad (7)$$

Once again, according to the properties (1) of disjunction:

$$\lfloor q_1 * r \rfloor \subseteq \lfloor q_1 * r \rfloor \cup \lfloor q_2 * r \rfloor = \lfloor (q_1 * r) \vee (q_2 * r) \rfloor = \lfloor (q_1 \vee q_2) * r \rfloor,$$

which together with (7) means that $(\sigma', \Sigma', \Delta') \in \lfloor (q_1 \vee q_2) * r \rfloor$. Overall we have shown that there exist Σ' and Δ' such that $\text{LP}^*(\Sigma, \Delta, \Sigma', \Delta')$ and $(\sigma', \Sigma', \Delta') \in \lfloor (q_1 \vee q_2) * r \rfloor$, which concludes the proof of (6).

We now prove the closure properties from Lemma 9.

Proof of FRAME. Let us define an auxiliary function:

$$\phi(X, r) \triangleq \{(p * r, C, q * r) \mid (p, C, q) \in X\}.$$

Then our goal is to prove that $\phi(\text{safe}_t, r) \subseteq \text{safe}_t$. Since $\text{safe}_t = \text{gfp } F_t$, we can do a proof by coinduction: to conclude that $\phi(\text{safe}_t, r) \subseteq \text{safe}_t$ holds, we demonstrate $\phi(\text{safe}_t, r) \subseteq F_t(\phi(\text{safe}_t, r))$.

Consider any $(p', C, q') \in \phi(\text{safe}_t, r)$. There necessarily are p, q, r such that $p' = p * r$, $q' = q * r$ and $(p, C, q) \in \text{safe}_t$. Let us assume that $C = \text{skip}$. Then $(p, C, q) \in \text{safe}_t$ implies that $p \Rightarrow q$. By Proposition 10, $p * r \Rightarrow q * r$, which implies $(p * r, \text{skip}, q * r) \in \text{safe}_t$ to hold.

Now let $C \neq \text{skip}$. Since $(p, C, q) \in \text{safe}_t$, by definition of the safety relation the following holds of every α, C' and any transition $C \xrightarrow{\alpha} C'$:

$$\exists p'. \alpha \Vdash_t \{p\}\{p'\} \wedge (p', C', q) \in \text{safe}_t.$$

By Proposition 10, $\alpha \Vdash_t \{p\}\{p'\}$ implies $\alpha \Vdash_t \{p * r\}\{p' * r\}$. Also, when $(p', C', q) \in \text{safe}_t$, it is the case that $(p' * r, C', q * r) \in \phi(\text{safe}_t, r)$. Thus, we have shown for every transition $C \xrightarrow{\alpha} C'$ that there exists $p'' = p' * r$ such that $\alpha \Vdash_t \{p * r\}\{p''\}$ and $(p'', C', q * r) \in \phi(\text{safe}_t, r)$:

$$\forall \alpha, C. C \xrightarrow{\alpha} C' \implies \exists p''. \alpha \Vdash_t \{p * r\}\{p''\} \wedge (p'', C', q * r) \in \phi(\text{safe}_t, r),$$

which is sufficient to conclude that $(p * r, C, q * r) \in F_t(\phi(\text{safe}_t, r))$. \square

Proof of CHOICE. Let us define an auxiliary function:

$$\phi(X, Y) \triangleq \{(p, C_1 + C_2, q) \mid (p, C_1, q) \in X \wedge (p, C_2, q) \in Y\}.$$

Then our goal is to prove that $\phi(\text{safe}_t, \text{safe}_t) \subseteq \text{safe}_t$. For convenience, we prove an equivalent inequality $\phi(\text{safe}_t, \text{safe}_t) \cup \text{safe}_t \subseteq \text{safe}_t$ instead.

Since $\text{safe}_t = \text{gfp } F_t$, we can do a proof by coinduction: to conclude that $\phi(\text{safe}_t, \text{safe}_t) \cup \text{safe}_t \subseteq \text{safe}_t$ holds, we demonstrate $\phi(\text{safe}_t, \text{safe}_t) \cup \text{safe}_t \subseteq F_t(\phi(\text{safe}_t, \text{safe}_t) \cup \text{safe}_t)$.

Let us consider $(p, C, q) \in \text{safe}_t$. Since $\text{safe}_t = \text{gfp } F_t$, we know that $\text{safe}_t = F_t(\text{safe}_t)$ holds. Then by monotonicity of F_t , $(p, C, q) \in F_t(\text{safe}_t) \subseteq F_t(\phi(\text{safe}_t, \text{safe}_t) \cup \text{safe}_t)$.

Now let us consider $(p, C, q) \in \phi(\text{safe}_t, \text{safe}_t)$. There necessarily are C_1 and C_2 such that $C = C_1 + C_2$, $(p, C_1, q) \in \text{safe}_t$, and $(p, C_2, q) \in \text{safe}_t$. For $(p, C_1 + C_2, q)$ to belong to $F_t(\phi(\text{safe}_t, \text{safe}_t) \cup \text{safe}_t)$, the following has to be proven for every transition $C_1 + C_2 \xrightarrow{\alpha} C'$:

$$\exists p'. \alpha \Vdash_t \{p\}\{p'\} \wedge (p', C', q) \in \phi(\text{safe}_t, \text{safe}_t) \cup \text{safe}_t. \quad (8)$$

According to the rules of the operational semantics (Figure 1), whenever $C_1 + C_2 \xrightarrow{\alpha} C'$, necessarily $\alpha = \text{id}$ and either $C' = C_1$ or $C' = C_2$. Let us assume that $C' = C_1$ (the other case is analogous). The action judgement $\text{id} \Vdash_t \{p\}\{p\}$ holds trivially. Knowing that $(p, C_1, q) \in \text{safe}_t$, it is easy to see that (8) can be satisfied by letting $p' = p$.

Consequently, $(p, C_1 + C_2, q) \in F_t(\phi(\text{safe}_t, \text{safe}_t) \cup \text{safe}_t)$, which concludes the proof. \square

Proof of DISJ. Let

$$\phi(X) \triangleq \{(p_1 \vee p_2, C, q_1 \vee q_2) \mid (p_1, C, q_1) \in X \wedge (p_2, C, q_2) \in X\}.$$

Then our goal is to prove that $\phi(\text{safe}_t) \subseteq \text{safe}_t$. Since $\text{safe}_t = \text{gfp } F_t$, we can do a proof by coinduction: to conclude that $\phi(\text{safe}_t) \subseteq \text{safe}_t$ holds, we demonstrate $\phi(\text{safe}_t) \subseteq F_t(\phi(\text{safe}_t))$.

Let us consider $(p, C, q) \in \phi(\text{safe}_t)$. Then there necessarily are p_1, q_1, p_2 and q_2 such that $p = p_1 \vee p_2$, $q = q_1 \vee q_2$, and $(p_1, C, q_1), (p_2, C, q_2) \in \text{safe}_t$. From the latter we get that for any α, C' and a transition $C \xrightarrow{\alpha} C'$ the following holds:

$$\begin{aligned} \exists p'_1. \alpha \Vdash_t \{p_1\}\{p'_1\} \wedge (p'_1, C', q_1) &\in \text{safe}_t; \\ \exists p'_2. \alpha \Vdash_t \{p_2\}\{p'_2\} \wedge (p'_2, C', q_2) &\in \text{safe}_t. \end{aligned}$$

Then it is the case that $(p'_1 \vee p'_2, C', q_1 \vee q_2) \in \phi(\text{safe}_t)$. Moreover, $\alpha \Vdash_t \{p_1 \vee p_2\}\{p'_1 \vee p'_2\}$ holds by Proposition 12. Thus, we have shown for every transition $C \xrightarrow{\alpha} C'$ that there exists $p' = p'_1 \vee p'_2$ such that $\alpha \Vdash_t \{p_1 \vee p_2\}\{p'\}$ and $(p', C', q_1 \vee q_2) \in \phi(\text{safe}_t)$:

$$\forall \alpha, C. C \xrightarrow{\alpha} C' \implies \exists p'. \alpha \Vdash_t \{p_1 \vee p_2\}\{p'\} \wedge (p', C', q_1 \vee q_2) \in \phi(\text{safe}_t).$$

which is sufficient to conclude that $(p_1 \vee p_2, C, q_1 \vee q_2) \in F_t(\phi(\text{safe}_t))$. \square

Proof of ITER. To do a proof by coinduction, we strengthen ITER property as follows:

$$\begin{aligned} \forall p, C. ((p, C, p) \in \text{safe}_t \implies (p, C^*, p) \in \text{safe}_t) \wedge \\ (\forall p_1, C_1. (p_1, C_1, p) \in \text{safe}_t \wedge (p, C, p) \in \text{safe}_t \implies (p_1, C_1 ; C^*, p) \in \text{safe}_t). \end{aligned} \quad (9)$$

Let us define auxilliary functions:

$$\begin{aligned} \phi(X) &\triangleq \{(p, C^*, p) \mid (p, C, p) \in X\} \\ \psi(X) &\triangleq \{(p_1, C_1 ; C_2^*, p_2) \mid (p_1, C_1, p_2), (p_2, C_2, p_2) \in X\}. \end{aligned}$$

Using them, we rewrite (9) as $\phi(\text{safe}_t) \cup \psi(\text{safe}_t) \subseteq \text{safe}_t$. Let $\xi = \{(p, \text{skip}, p)\}$. It is easy to see that $\phi(\text{safe}_t) \cup \psi(\text{safe}_t) \cup \xi \subseteq \text{safe}_t$ is also an equivalent reformulation of (9), since $\xi \subseteq \text{safe}_t$ always holds.

Since $\text{safe}_t = \text{gfp } F_t$, we can do a proof by coinduction: to conclude that $\phi(\text{safe}_t) \cup \psi(\text{safe}_t) \cup \xi \subseteq \text{safe}_t$, we demonstrate $\phi(\text{safe}_t) \cup \psi(\text{safe}_t) \cup \xi \subseteq F_t(\phi(\text{safe}_t) \cup \psi(\text{safe}_t) \cup \xi)$.

Consider any $(p, C, q) \in \xi$. Necessarily, $C = \text{skip}$ and $q = p$. Note that $p \Rightarrow p$ always holds, which by definition of F_t is sufficient for $(p, \text{skip}, p) \in F_t(\phi(\text{safe}_t) \cup \psi(\text{safe}_t) \cup \xi)$. Thus, $(p, C, q) \in F_t(\phi(\text{safe}_t) \cup \psi(\text{safe}_t) \cup \xi)$.

Consider any $(p, C', q) \in \phi(\text{safe}_t)$. Necessarily, $p = q$ and there exists a sequential command C such that $C' = C^*$ and $(p, C, p) \in \text{safe}_t$. We need to show that $(p, C^*, p) \in F_t(\phi(\text{safe}_t) \cup \psi(\text{safe}_t) \cup \xi)$. For the latter to hold, by definition of F_t it is sufficient that for every α, C'' and a transition $C \xrightarrow{\alpha} C''$ the following be true:

$$\exists p''. \alpha \Vdash_t \{p\}\{p''\} \wedge (p'', C'', p) \in \phi(\text{safe}_t) \cup \psi(\text{safe}_t) \cup \xi. \quad (10)$$

According to the operational semantics in Figure 1, when there is a transition $C \xrightarrow{\alpha} C''$, necessarily $\alpha = \text{id}$ and either $C'' = \text{skip}$ or $C'' = C ; C^*$. Let us assume that $C'' = \text{skip}$. Since both $(p, \text{skip}, p) \in \xi$ and $\alpha \Vdash_t \{p\}\{p\}$ always hold, it is easy to see that letting $p'' = p$ satisfies (10). Now let us turn to the case when $C'' = C ; C^*$. Note that $(p, C ; C^*, p) \in \psi(\text{safe}_t)$ holds by definition of ψ . Thus, by letting $p'' = p$ we satisfy (10).

Consider $(p_1, C_0, p_2) \in \psi(\text{safe}_t)$. Necessarily, there exist C_1 and C_2 such that:

$$C_0 = C_1 ; C_2^* \wedge (p_1, C_1, p_2) \in \text{safe}_t \wedge (p_2, C_2, p_2) \in \text{safe}_t. \quad (11)$$

We need to show that $(p_1, C_1 ; C_2^*, p_2) \in F_t(\phi(\text{safe}_t) \cup \psi(\text{safe}_t) \cup \xi)$. For the latter to hold, we need to prove the following for every α, C' and a transition $C_1 ; C_2^* \xrightarrow{\alpha} C'$:

$$\exists p'. \alpha \Vdash_t \{p_1\}\{p'\} \wedge (p', C', p_2) \in \phi(\text{safe}_t) \cup \psi(\text{safe}_t) \cup \xi. \quad (12)$$

According to the operational semantics in Figure 1, when there is a transition $C_1 ; C_2^* \xrightarrow{\alpha} C'$, either of the following is true:

- there are C'_1 and a transition $C_1 \xrightarrow{C} C'_1$ such that $C' = C'_1 ; C_2^*$;
- $C_1 = \text{skip}$, $C' = C_2$ and $\alpha = \text{id}$.

Let us assume that the former is the case. From (11) we know that $(p_1, C_1, p_2) \in \text{safe}_t$, so by definition of the safety relation we get that:

$$\exists p'_1. \alpha \Vdash_t \{p_1\}\{p'_1\} \wedge (p'_1, C'_1, p_2) \in \text{safe}_t.$$

Consequently, $(p'_1, C'_1 ; C_2^*, p_2) \in \psi(\text{safe}_t)$. Thus, by letting $p' = p'_1$ we can satisfy (12).

Now let $C_1 = \text{skip}$ and $\alpha = \text{id}$. From (11) we know that $(p_1, \text{skip}, p_2) \in \text{safe}_t$, meaning that necessarily $p_1 \Rightarrow p_2$. It is easy to see that $p_1 \Rightarrow p_2$ holds if and only if so does $\text{id} \Vdash_t \{p_1\}\{p_2\}$. Knowing that $(p_2, C_2, p_2) \in \text{safe}_t$, we can satisfy (12) by letting $p' = p_2$. \square

Proof of SEQ. Let

$$\phi(X, q') \triangleq \{(p, C_1 ; C_2, q) \mid (p, C_1, q') \in X \wedge (q', C_2, q) \in X\}.$$

Then our goal is to prove that $\phi(\text{safe}_t, q') \subseteq \text{safe}_t$. For convenience, we prove an equivalent inequality $\phi(\text{safe}_t, q') \cup \text{safe}_t \subseteq \text{safe}_t$ instead.

Since $\text{safe}_t = \text{gfp } F_t$, we can do a proof by coinduction: to conclude that $\phi(\text{safe}_t, q') \cup \text{safe}_t \subseteq \text{safe}_t$ holds, we demonstrate $\phi(\text{safe}_t, q') \cup \text{safe}_t \subseteq F_t(\phi(\text{safe}_t, q') \cup \text{safe}_t)$.

Let us consider any $(p, C, q) \in \text{safe}_t$. Since $\text{safe}_t = \text{gfp } F_t$, we know that $\text{safe}_t = F_t(\text{safe}_t)$ holds. Then by monotonicity of F_t , $(p, C, q) \in F_t(\text{safe}_t) \subseteq F_t(\phi(\text{safe}_t, \text{safe}_t) \cup \text{safe}_t)$.

Now let us consider any $(p, C, q) \in \phi(\text{safe}_t, q')$. There necessarily are C_1 and C_2 such that:

$$C = C_1 ; C_2 \wedge (p, C_1, q') \in \text{safe}_t \wedge (q', C_2, q) \in \text{safe}_t. \quad (13)$$

For $(p, C_1 ; C_2, q)$ to belong to $F_t(\phi(\text{safe}_t, \text{safe}_t) \cup \text{safe}_t)$, the following has to be the case for every transition $C_1 ; C_2 \xrightarrow{\alpha} C'$:

$$\exists p'. \alpha \Vdash_t \{p\}\{p'\} \wedge (p', C', q) \in \phi(\text{safe}_t, q') \cup \text{safe}_t. \quad (14)$$

According to the rules of the operational semantics (Figure 1), when there is a transition $C_1 ; C_2 \xrightarrow{\alpha} C'$, either of the following is true:

- there exists C'_1 such that $C' = C'_1$; C_2 and $C_1 \xrightarrow{\alpha} C'_1$; or
- $C_1 = \text{skip}$, $\alpha = \text{id}$ and $C' = C_2$.

Let us assume that the former is the case. From (13) we know that $(p, C_1, q') \in \text{safe}_t$, which means that the following holds of $C_1 \xrightarrow{\alpha} C'_1$:

$$\exists p''. \alpha \Vdash_t \{p\}\{p''\} \wedge (p'', \alpha'_1, q') \in \text{safe}_t.$$

When $(p'', \alpha'_1, q') \in \text{safe}_t$ and $(q', \alpha_2, q) \in \text{safe}_t$, it is the case that $(p'', \alpha'_1 ; \alpha_2, q) \in \phi(\text{safe}_t, q')$. Thus, by letting $p' = p''$ we satisfy (14).

We now consider the case when $C_1 = \text{skip}$, $\alpha = \text{id}$ and $C' = C_2$. From (13) we know that $(p, \text{skip}, q') \in \text{safe}_t$, meaning that $p \Rightarrow q'$, or equivalently $\text{id} \Vdash_t \{p\}\{q'\}$. We also know from (13) that $(q', C_2, q) \in \text{safe}_t$. Thus, (14) can be satisfied by letting $p' = q'$. \square

Proof of CONSEQ. Let us first show that $\text{safe}_t(p', C, q)$ holds, when so do $p' \Rightarrow p$ and $\text{safe}_t(p, C, q)$. When $C = \text{skip}$, $\text{safe}_t(p, C, q)$ gives us that $p \Rightarrow q$. It is easy to see that $p' \Rightarrow p$ and $p \Rightarrow q$ together imply $p' \Rightarrow q$, which is sufficient to conclude that $\text{safe}_t(p', C, q)$ holds. Let us assume that $C \neq \text{skip}$. From $\text{safe}_t(p, C, q)$ we get that the following holds of every transition $C \xrightarrow{\alpha} C'$:

$$\exists p''. \alpha \Vdash_t \{p\}\{p''\} \wedge \text{safe}_t(p'', C', q)$$

However, by applying Proposition 11 about Consequence property of axiom judgements to $p' \Rightarrow p$ and $\alpha \Vdash_t \{p\}\{p''\}$ we get that $\alpha \Vdash_t \{p'\}\{p''\}$. Together with the formula above, it allows us to conclude that $\text{safe}_t(p', C, q)$ holds.

Now let us prove that $\text{safe}_t(p, C, q')$ holds, when so do $q \Rightarrow q'$ and $\text{safe}_t(p, C, q)$. We define an auxilliary function:

$$\phi(X, q) \triangleq \{(p, C, q') \mid (p, C, q) \in X \wedge q \Rightarrow q'\}.$$

Our goal is to prove that $\phi(\text{safe}_t, q) \subseteq \text{safe}_t$. Since $\text{safe}_t = \text{gfp } F_t$, we can do a proof by coinduction: to conclude that $\phi(\text{safe}_t, q) \subseteq \text{safe}_t$ holds, we demonstrate $\phi(\text{safe}_t, q) \subseteq F_t(\phi(\text{safe}_t, q))$.

Let us consider any $(p, C, q') \in \phi(\text{safe}_t, q)$. Necessarily, $(p, C, q) \in \text{safe}_t$ and $q \Rightarrow q'$. When $C = \text{skip}$, we need to show that $p \Rightarrow q'$. Since $(p, \text{skip}, q) \in \text{safe}_t$, it is the case that $p \Rightarrow q$. It is easy to see that $p \Rightarrow q$ and $q \Rightarrow q'$ together imply $p \Rightarrow q'$, which is sufficient to conclude that $(p, C, q') \in F_t(\phi(\text{safe}_t, q))$.

Now consider the case when $C \neq \text{skip}$. Since $(p, \alpha, q) \in \text{safe}_t$, by definition of the safety relation the following holds of every α, C' and a transition $C \xrightarrow{\alpha} C'$:

$$\exists p''. \alpha \Vdash_t \{p\}\{p''\} \wedge (p'', C', q) \in \text{safe}_t$$

Knowing that $q \Rightarrow q'$ and $(p'', C', q) \in \text{safe}_t$, it is easy to see that $(p'', C', q') \in \phi(\text{safe}_t, q)$. Thus, we have shown that:

$$\forall \alpha, C'. C \xrightarrow{\alpha} C' \implies \exists p''. \alpha \Vdash_t \{p\}\{p''\} \wedge (p'', C', q') \in \phi(\text{safe}_t, q),$$

which is sufficient for $(p, C, q') \in F_t(\phi(\text{safe}_t, q))$ to hold. \square

C Proof of Lemma 4

Lemma. $\forall t, \mathcal{P}, C, \mathcal{Q}. \vdash_t \{\mathcal{P}\} C \{\mathcal{Q}\} \implies \forall i. \text{safe}_t(\llbracket \mathcal{P} \rrbracket_i, C, \llbracket \mathcal{Q} \rrbracket_i).$

We prove Lemma 4 by rule induction. For that we choose arbitrary thread identifier t and demonstrate that $\forall i. \text{safe}_t(\llbracket \mathcal{P} \rrbracket_i, C, \llbracket \mathcal{Q} \rrbracket_i)$ is closed under the proof rules from Figure 3. The cases of CHOICE, ITER, SEQ, CONSEQ, FRAME and DISJ rules are straightforward: they trivially follow from Lemma 9 after using the properties of $\llbracket - \rrbracket_i$ from Figure 2. The EX rule uses the fact that Val, which is the range of i , is finite, which makes possible proving it just like the DISJ rule.

It remains to consider the PRIM rule to conclude Lemma 4. Let us assume that $\forall i'. \alpha \Vdash_t \{\llbracket \mathcal{P} \rrbracket_{i'}\} \{\llbracket \mathcal{Q} \rrbracket_{i'}\}$ holds. We need to demonstrate that so does $\forall i. \text{safe}_t(\llbracket \mathcal{P} \rrbracket_i, \alpha, \llbracket \mathcal{Q} \rrbracket_i)$. To conclude that the latter holds, according Definition 3 we need to prove the following for every i :

$$\forall C', \alpha'. \alpha \xrightarrow{\alpha'} C' \implies \exists p'. \alpha' \Vdash_t \{\llbracket \mathcal{P} \rrbracket_i\} \{p'\} \wedge \text{safe}_t(p', C', \llbracket \mathcal{Q} \rrbracket_i) \quad (15)$$

According to the operational semantics from Figure 1, the only transition from a command α is $\alpha \xrightarrow{\alpha} \text{skip}$. Thus, in the formula above $\alpha' = \alpha$ and $C' = \text{skip}$. Note that $\text{safe}_t(\llbracket \mathcal{Q} \rrbracket_i, \text{skip}, \llbracket \mathcal{Q} \rrbracket_i)$ holds trivially. Additionally, by our assumption, $\alpha \Vdash_t \{\llbracket \mathcal{P} \rrbracket_{i'}\} \{\llbracket \mathcal{Q} \rrbracket_{i'}\}$ holds for any i' . Consequently, it holds for $i' = i$. We conclude that by letting $p' = \llbracket \mathcal{Q} \rrbracket_i$ we satisfy (15). \square

D Proof of Theorem 7

We further refer to the assumptions of Theorem 7 as a relation $\text{safelib}(\ell, \mathcal{L}, \mathcal{P}, \mathcal{Q})$ defined as follows.

Definition 13 *Given a concrete library ℓ , an abstract library \mathcal{L} and $\mathcal{P}, \mathcal{Q} : \text{ThreadID} \times \text{APCom} \rightarrow \text{VAssn}$, we say that a relation $\text{safelib}(\ell, \mathcal{L}, \mathcal{P}, \mathcal{Q})$ holds if and only if the following requirements are met:*

1. $\text{dom}(\ell) = \text{dom}(\mathcal{L});$
2. $\forall i, t, A, \sigma, \Sigma, \Delta, r. (\sigma, \Sigma, \Delta) \in \llbracket \mathcal{P}(t, A) \rrbracket_i * r \implies \Delta(t) = \text{todo}(A);$
3. $\forall i, t, A, \sigma, \Sigma, \Delta, r. (\sigma, \Sigma, \Delta) \in \llbracket \mathcal{Q}(t, A) \rrbracket_i * r \implies \Delta(t) = \text{done}(A);$
4. $\forall i, t, A, A', r, \Delta. ((\sigma, \Sigma, \Delta[t : \text{todo}(A)]) \in \llbracket \mathcal{P}(t, A) \rrbracket_i * r \iff (\sigma, \Sigma, \Delta[t : \text{done}(A')]) \in \llbracket \mathcal{Q}(t, A') \rrbracket_i * r);$
5. $\forall m, a, r, t. m \in \text{dom}(\ell) \wedge a, r \in \text{Val} \wedge t \in \text{ThreadID} \implies \vdash_t \{\mathcal{P}(t, \mathcal{L}(m, a, r))\} \ell(m, a, r) \{\mathcal{Q}(t, \mathcal{L}(m, a, r))\};$

To strengthen the statement of Theorem 7 as necessary for its proof, we define an auxilliary relation, a *thread pool invariant*. With this relation we establish a correspondence between the information about LP in a thread t from a given view v_t and sequential commands in a thread t of a concrete thread pool τ and abstract thread pool \mathcal{T} .

Definition 14 Given a concrete library ℓ , an abstract library \mathcal{L} , predicates $\mathcal{P}, \mathcal{Q} : \text{ThreadID} \times \text{APCom} \rightarrow \text{VAssn}$, a concrete thread pool τ , an abstract thread pool \mathcal{T} , a view v_t and an interpretation of logical variables \mathbf{i} , we say that a thread pool invariant $\text{inv}_t(\mathbf{i}, \tau, \mathcal{T}, v_t, \Delta)$ holds in a thread t if and only if the following requirements are met:

- if $\tau(t) = \text{idle}$, then $\mathcal{T}(t) = \text{idle}$ and $v_t \Rightarrow \llbracket \mathcal{Q}(t, _) \rrbracket_{\mathbf{i}}$, or
- there exist C, r, m, a such that $\tau(t) = (C, r)$ and the following holds:

$$\text{safe}_t(v_t, C, \mathcal{Q}(t, \mathcal{L}(m, a, r))) \wedge ((\Delta(t) = \text{todo}(\mathcal{L}(m, a, r)) \wedge \mathcal{T}(t) = (\mathcal{L}(m, a, r), r)) \vee (\Delta(t) = \text{done}(\mathcal{L}(m, a, r)) \wedge \mathcal{T}(t) = (\text{skip}, r))).$$

Finally, analogously to Definition 5, we write down formally a definition of the set of histories of abstract libraries.

Definition 15 We define $\mathcal{H}_n[\mathcal{L}, \mathcal{T}, \Sigma]$ as a set of histories such that $\mathcal{H}_0[\mathcal{L}, \mathcal{T}, \Sigma] \triangleq \{\varepsilon\}$ and:

$$\begin{aligned} \mathcal{H}_n[\mathcal{L}, \mathcal{T}, \Sigma] \triangleq & \{((t, \text{call } m(a)) :: h) \mid m \in \text{dom}(\mathcal{L}) \wedge \mathcal{T}(t) = \text{idle} \wedge \\ & \exists r. h \in \mathcal{H}_{n-1}[\mathcal{L}, \mathcal{T}[t : (\mathcal{L}(m, a, r), r)], \Sigma]\} \\ & \cup \{h \mid \exists t, A, \Sigma', r. \mathcal{T}(t) = (A, r) \wedge \Sigma' \in \llbracket A \rrbracket_t(\Sigma) \wedge \\ & h \in \mathcal{H}_{n-1}[\mathcal{L}, \mathcal{T}[t : (\text{skip}, r)], \Sigma']\} \\ & \cup \{((t, \text{ret } m(r)) :: h) \mid m \in \text{dom}(\mathcal{L}) \wedge \mathcal{T}(t) = (\text{skip}, r) \wedge \\ & h \in \mathcal{H}_{n-1}[\mathcal{L}, \mathcal{T}[t : \text{idle}], \Sigma]\} \end{aligned}$$

We let $\mathcal{H}[\ell, \sigma] = \bigcup_{n \geq 0} \mathcal{H}_n[\mathcal{L}, (\lambda t. \text{idle}), \Sigma]$ denote the set of all possible histories of a library \mathcal{L} that start from a state Σ .

We are now ready to prove Theorem 7.

Proof. Let us consider any $\ell, \mathcal{L}, \mathcal{P}, \mathcal{Q}$ such that $\text{safelib}(\ell, \mathcal{L}, \mathcal{P}, \mathcal{Q})$ holds. Let us explain how we strengthen the statement of the theorem in this proof. We prove that $\forall n. \phi(n)$ holds with $\phi(n)$ formulated as follows:

$$\begin{aligned} \phi(n) = & \forall \mathbf{i}, \sigma, \Sigma, \Delta, \tau, \mathcal{T}. (\exists v_1, \dots, v_N. (\forall k. \text{inv}_k(\mathbf{i}, \tau, \mathcal{T}, v_k, \Delta)) \wedge \\ & (\sigma, \Sigma, \Delta) \in \llbracket \otimes_{t \in \text{ThreadID}} v_t \rrbracket) \implies \mathcal{H}_n[\ell, \tau, \sigma] \subseteq \mathcal{H}[\mathcal{L}, \mathcal{T}, \Sigma]. \end{aligned} \quad (16)$$

Note that according to the semantics of the assertion language Assn (Figure 2):

$$\llbracket \otimes_{k \in \text{ThreadID}} (\exists A. \mathcal{Q}(k, A)) \rrbracket_{\mathbf{i}} = \otimes_{k \in \text{ThreadID}} \llbracket (\exists A. \mathcal{Q}(k, A)) \rrbracket_{\mathbf{i}}.$$

With that in mind, it is easy to see that letting $v_k = \llbracket (\exists A. \mathcal{Q}(k, A)) \rrbracket_{\mathbf{i}}$ for all $k \in \text{ThreadID}$, $\tau = (\lambda t. \text{idle})$ and $\mathcal{T} = (\lambda t. \text{idle})$ in (16) yields the formula:

$$\begin{aligned} (\forall \mathbf{i}, \sigma, \Sigma, \Delta. (\sigma, \Sigma, \Delta) \in \llbracket \otimes_{k \in \text{ThreadID}} (\exists A. \mathcal{Q}(k, A)) \rrbracket_{\mathbf{i}}) \implies \\ \bigcup_{n \geq 0} \mathcal{H}_n[\ell, \lambda t. \text{idle}, \sigma] \subseteq \mathcal{H}[\mathcal{L}, \lambda t. \text{idle}, \Sigma], \end{aligned}$$

which coincides with the statement of the theorem.

We prove $\forall n. \phi(n)$ by induction on n . Let us take any $\mathbf{i}, \sigma, \Sigma, \Delta, \tau$ and \mathcal{T} , and consider v_1, \dots, v_N such that the premisses of $\phi(n)$ hold:

$$(\forall k. \text{inv}_k(\mathbf{i}, \tau, \mathcal{T}, v_k, \Delta)) \wedge (\sigma, \Sigma, \Delta) \in \llbracket \otimes_{k \in \text{ThreadID}} v_k \rrbracket \quad (17)$$

We need to demonstrate that every history h of the concrete library ℓ from the set $\mathcal{H}_n[\ell, \tau, \sigma]$ is also a history of the abstract library \mathcal{L} : $h \in \mathcal{H}[\mathcal{L}, \mathcal{T}, \Sigma]$.

By Definition 5 of $\mathcal{H}_n[\ell, \tau, \sigma]$, if $n = 0$, then h is an empty history that is trivially present in $\mathcal{H}[\mathcal{L}, \mathcal{T}, \Sigma]$. Let us now consider $n > 0$ and assume that $\phi(n - 1)$ holds. By definition of $\mathcal{H}_n[\ell, \tau, \sigma]$, h corresponds to one of the three events in a thread t : a call of an arbitrary method m with an argument a in a thread t , a return from a method m with a return value r or a transition in a thread t . We consider each case separately.

Case #1. There is a history h' , a thread t , a method $m \in \text{dom}(\ell)$, its argument a and a return value r such that $h = (t, \text{call } m(a)) :: h'$, $\tau(t) = \text{idle}$ and $h' \in \mathcal{H}_{n-1}[\ell, \tau[t : (\ell(m, a, r), r), \sigma]]$. By Definition 15, to conclude that $h = (t, \text{call } m(a)) :: h' \in \mathcal{H}[\mathcal{L}, \mathcal{T}, \Sigma]$ it is necessary to show that $\mathcal{T}(t) = \text{idle}$ and $h' \in \mathcal{H}[\mathcal{L}, \mathcal{T}[t : (\mathcal{L}(m, a, r), r)], \Sigma]$, which we further do in the proof of Case #1.

According to (17), $\text{inv}_t(\mathbf{i}, \tau, \mathcal{T}, v_t, \Delta)$ and $(\sigma, \Sigma, \Delta) \in \lfloor \otimes_{k \in \text{ThreadID}} v_k \rfloor$ hold. Then necessarily $\mathcal{T}(t) = \text{idle}$ and $v_t \Rightarrow \llbracket Q(t, _) \rrbracket_{\mathbf{i}}$, which corresponds to the only case when $\tau(t) = \text{idle}$ in the thread pool invariant. By Definition 5 of $v_t \Rightarrow \llbracket Q(t, _) \rrbracket_{\mathbf{i}}$, $(\sigma, \Sigma, \Delta) \in \lfloor \otimes_{k \in \text{ThreadID}} v_k \rfloor$ implies $(\sigma, \Sigma, \Delta) \in \llbracket \llbracket Q(t, _) \rrbracket_{\mathbf{i}} * \otimes_{k \in \text{ThreadID} \setminus \{t\}} v_k \rfloor$. From requirements to predicates \mathcal{P} and \mathcal{Q} in $\text{safelib}(\ell, \mathcal{L}, \mathcal{P}, \mathcal{Q})$ we obtain that the following holds of (σ, Σ, Δ) :

- $\Delta(t) = \text{done}(_)$, and
- $(\sigma, \Sigma, \Delta[t : \text{todo}(\mathcal{L}(m, a, r))]) \in \llbracket \llbracket \mathcal{P}(t, \mathcal{L}(m, a, r)) \rrbracket_{\mathbf{i}} * \otimes_{k \in \text{ThreadID} \setminus \{t\}} v_k \rfloor$.

Let $v'_t = \llbracket \mathcal{P}(t, \mathcal{L}(m, a, r)) \rrbracket_{\mathbf{i}}$ and $v'_k = v_k$ for all $k \neq t$. Obviously, $(\sigma, \Sigma, \Delta[t : \text{todo}(\mathcal{L}(m, a, r))]) \in \lfloor \otimes_k v'_k \rfloor$.

Also, by Lemma 4, $\text{safe}_t(\llbracket \mathcal{P}(t, \mathcal{L}(m, a, r)) \rrbracket_{\mathbf{i}}, \ell(m, a, r), \llbracket Q(t, \mathcal{L}(m, a, r)) \rrbracket_{\mathbf{i}})$ holds. This allows us to conclude that in a thread t a thread pool invariant $\text{inv}_t(\mathbf{i}, \tau[t : (\ell(m, a, r), r)], \mathcal{T}[t : (\mathcal{L}(m, a, r), r)], v'_t, \Delta[t : \mathcal{L}(m, a, r)])$ holds. Moreover, according to (17), thread pool invariants hold in all other threads as well.

We have shown that there exist v'_1, \dots, v'_N such that:

$$(\forall k. \text{inv}_k(\mathbf{i}, \tau[t : (\ell(m, a, r), r)], \mathcal{T}[t : (\mathcal{L}(m, a, r), r)], v'_k, \Delta[t : \mathcal{L}(m, a, r)])) \wedge (\sigma, \Sigma, \Delta[t : \text{todo}(\mathcal{L}(m, a, r))]) \in \lfloor \otimes_{k \in \text{ThreadID}} v'_k \rfloor,$$

which by the induction hypothesis $\phi(n - 1)$ implies that $h' \in \mathcal{H}_{n-1}[\ell, \tau[t : (\ell(m, a, r), r)], \sigma] \subseteq \mathcal{H}[\mathcal{L}, \mathcal{T}[t : (\mathcal{L}(m, a, r), r)], \Sigma]$. We have also show that $\mathcal{T}(t) = \text{idle}$. By Definition 15, $h = (t, \text{call } m(a)) :: h' \in \mathcal{H}[\mathcal{L}, \mathcal{T}, \Sigma]$, which concludes the proof of Case #1.

Case #2. There is a history h' , a thread t , a method $m \in \text{dom}(\ell)$, its argument a and a return value r such that $h = (t, \text{ret } m(r)) :: h'$, $\tau(t) = (\text{skip}, r)$ and $h' \in \mathcal{H}_{n-1}[\ell, \tau[t : \text{idle}], \sigma]$. By Definition 15, to conclude that $h = (t, \text{ret } m(r)) :: h' \in \mathcal{H}[\mathcal{L}, \mathcal{T}, \Sigma]$ it is necessary to show that $\mathcal{T}(t) = (\text{skip}, r)$ and $h' \in \mathcal{H}[\mathcal{L}, \mathcal{T}[t : \text{idle}], \Sigma]$, which we further do in this proof of Case #2.

According to (17), a thread invariant $\text{inv}_t(\mathbf{i}, \tau, \mathcal{T}, v_t, \Delta)$ holds. Then the following is true:

$$\text{safe}_t(v_t, \text{skip}, \llbracket Q(t, \mathcal{L}(m, a, r)) \rrbracket_{\mathbf{i}}) \wedge ((\Delta(t) = \text{todo}(\mathcal{L}(m, a, r)) \wedge \mathcal{T}(t) = (\mathcal{L}(m, a, r), r)) \vee (\Delta(t) = \text{done}(\mathcal{L}(m, a, r)) \wedge \mathcal{T}(t) = (\text{skip}, r))). \quad (18)$$

By Definition 3 of $\text{safe}_t(v_t, \text{skip}, \llbracket \mathcal{Q}(t, \mathcal{L}(m, a, r)) \rrbracket_i)$, $v_t \Rightarrow \llbracket \mathcal{Q}(t, \mathcal{L}(m, a, r)) \rrbracket_i$ holds. Consequently, by Definition 5:

$$(\sigma, \Sigma, \Delta) \in \llbracket v_t * \otimes_{k \in \text{ThreadID} \setminus \{t\}} v_k \rrbracket \subseteq \llbracket \llbracket \mathcal{Q}(t, \mathcal{L}(m, a, r)) \rrbracket_i * \otimes_{k \in \text{ThreadID} \setminus \{t\}} v_k \rrbracket.$$

From the third requirement to \mathcal{Q} in Definition 13:

$$(\sigma, \Sigma, \Delta) \in \llbracket \llbracket \mathcal{Q}(t, \mathcal{L}(m, a, r)) \rrbracket_i * \otimes_k v_k \rrbracket \implies \Delta(t) = \text{done}(\mathcal{L}(m, a, r)).$$

Consequently, from (18) we get that $\Delta(t) = \text{done}(\mathcal{L}(m, a, r))$ and $\mathcal{T}(t) = (\text{skip}, r)$.

Let $v'_t = \llbracket \mathcal{Q}(t, \mathcal{L}(m, a, r)) \rrbracket_i$ and $v'_k = v_k$ for $k \neq t$. It is easy to see that $\text{inv}_t(\mathbf{i}, \tau[t : \text{idle}], \tau[t : \text{idle}], \mathcal{T}[t : \text{idle}], v'_t, \Delta)$ holds trivially by Definition 14. Moreover, according to (17), thread pool invariants hold in other threads as well.

We have shown that there exist v'_1, \dots, v'_N such that:

$$(\forall k. \text{inv}_k(\mathbf{i}, \tau[t : \text{idle}], \mathcal{T}[t : \text{idle}], v'_k, \Delta)) \wedge (\sigma, \Sigma, \Delta) \in \llbracket \otimes_{k \in \text{ThreadID}} v'_k \rrbracket,$$

which by the induction hypothesis $\phi(n-1)$ implies that $h' \in \mathcal{H}_{n-1}[\ell, \tau[t : \text{idle}], \sigma] \subseteq \mathcal{H}[\mathcal{L}, \mathcal{T}[t : \text{idle}], \Sigma]$. We have also show that $\mathcal{T}(t) = (\text{skip}, r)$. By Definition 15, $h = (t, \text{ret } m(r)) :: h' \in \mathcal{H}[\mathcal{L}, \mathcal{T}, \Sigma]$, which concludes the proof of Case #2.

Case #3. There is a thread t , sequential commands C and C' , a primitive command α , concrete states σ and σ' and a return value r such that $\tau(t) = (C, r)$, $\langle C, \sigma \rangle \xrightarrow{t, \alpha} \langle C', \sigma' \rangle$ and $h \in \mathcal{H}_{n-1}[\ell, \tau[t : (C', r)], \sigma']$.

According to (17), $\text{inv}_t(\mathbf{i}, \tau, \mathcal{T}, v_t, \Delta)$ holds. Consequently, there exist a method m with its argument a such that $C = \mathcal{L}(m, a, r)$ and:

$$\begin{aligned} & \text{safe}_t(v_t, C, \llbracket \mathcal{Q}(t, \mathcal{L}(m, a, r)) \rrbracket_i) \wedge \\ & ((\Delta(t) = \text{todo}(\mathcal{L}(m, a, r)) \wedge \mathcal{T}(t) = (\mathcal{L}(m, a, r), r)) \vee \\ & (\Delta(t) = \text{done}(\mathcal{L}(m, a, r)) \wedge \mathcal{T}(t) = (\text{skip}, r))). \end{aligned} \quad (19)$$

It is easy to see that whenever there is a transition $\langle C, \sigma \rangle \xrightarrow{t, \alpha} \langle C', \sigma' \rangle$, there also is a stateless transition $C \xrightarrow{\alpha} C'$. By Definition 3 of $\text{safe}_t(v_t, C, \llbracket \mathcal{Q}(t, \mathcal{L}(m, a, r)) \rrbracket_i)$, if $C \xrightarrow{\alpha} C'$, then there exists a view v'_t such that $\alpha \Vdash_t \{v_t\} \{v'_t\}$ and $\text{safe}_t(v'_t, C', \llbracket \mathcal{Q}(t, \mathcal{L}(m, a, r)) \rrbracket_i)$.

Let $v'_k = v_k$ for any $k \neq t$. By Definition 2 of the action judgement $\alpha \Vdash_t \{v_t\} \{v'_t\}$, for $(\sigma, \Sigma, \Delta) \in \llbracket v_t * \otimes_{k \neq t} v_k \rrbracket$ and any $\sigma' \in \llbracket \alpha \rrbracket_t(\sigma)$, there exist Σ', Δ' such that:

$$\text{LP}^*(\Sigma, \Delta, \Sigma', \Delta') \wedge (\sigma', \Sigma', \Delta') \in \llbracket v'_t * \otimes_{k \neq t} v_k \rrbracket. \quad (20)$$

Let us assume that $\Delta = \Delta'$. Note that $\text{safe}_t(v'_t, C', \llbracket \mathcal{Q}(t, \mathcal{L}(m, a, r)) \rrbracket_i)$ holds, and according to (19) the following holds too:

$$\begin{aligned} & (\Delta(t) = \text{todo}(\mathcal{L}(m, a, r)) \wedge \mathcal{T}(t) = (\mathcal{L}(m, a, r), r)) \vee \\ & (\Delta(t) = \text{done}(\mathcal{L}(m, a, r)) \wedge \mathcal{T}(t) = (\text{skip}, r)) \end{aligned}$$

Thus, it is easy to see that $\text{inv}_t(\mathbf{i}, \tau[t : (C', r)], \mathcal{T}, v'_t, \Delta)$ holds. Combining this observation with (20), we conclude that we have demonstrated existence of v'_1, \dots, v'_N such that:

$$(\forall k. \text{inv}_k(\mathbf{i}, \tau[t : (C', r)], \mathcal{T}, v'_k, \Delta)) \wedge (\sigma', \Sigma', \Delta) \in \llbracket \otimes_{k \in \text{ThreadID}} v'_k \rrbracket,$$

which by the induction hypothesis $\phi(n-1)$ implies that $h \in \mathcal{H}_{n-1}[\ell, \tau[t : (C', r)], \sigma'] \subseteq \mathcal{H}[\mathcal{L}, \mathcal{T}, \Sigma]$. This concludes the proof of the case when $\Delta = \Delta'$.

We now return to the case when $\Delta \neq \Delta'$. According to (20), $\text{LP}^*(\Sigma, \Delta, \Sigma', \Delta')$ holds, meaning that linearization points of one or more threads have been passed. Without loss of generality, we assume the case of exactly one linearization point, i.e. that $\text{LP}(\Sigma, \Delta, \Sigma', \Delta')$ holds. Consequently, according to Definition 2 there exist t' and A' such that:

$$\Sigma' \in \llbracket A' \rrbracket_{t'}(\Sigma) \wedge \Delta(t') = \text{todo}(A') \wedge \Delta' = \Delta[t' : \text{done}(A')] \quad (21)$$

Let us consider the thread pool invariant $\text{inv}_{t'}(\mathbf{i}, \tau, \mathcal{T}, v_{t'}, \Delta)$, which holds according to (17). We show that $\tau(t') \neq \text{idle}$. From (21) we know that $\Delta(t') = \text{todo}(A')$. Since the third requirement to \mathcal{Q} in Definition 13 requires that $\Delta(t) = \text{done}(\mathcal{L}(m, a, r))$ hold, by Definition 14 it can only be the case that there exist C'', m', a', r' such that $\tau(t') = (C'', r')$ and the following is true:

$$\begin{aligned} & \text{safe}_{t'}(v_t, C'', \llbracket \mathcal{Q}(t', \mathcal{L}(m', a', r')) \rrbracket_{\mathbf{i}}) \wedge \\ & ((\Delta(t') = \text{todo}(\mathcal{L}(m', a', r')) \wedge \mathcal{T}(t') = (\mathcal{L}(m', a', r'), r')) \vee \\ & (\Delta(t') = \text{done}(\mathcal{L}(m', a', r')) \wedge \mathcal{T}(t') = (\text{skip}, r'))). \end{aligned} \quad (22)$$

From formula (21) we know that $A' = \mathcal{L}(m', a', r')$. Consequently, $\Delta'(t') = \text{done}(\mathcal{L}(m', a', r')) \wedge \mathcal{T}[t' : (\text{skip}, r')](t') = (\text{skip}, r')$ holds, which allows us to conclude the thread pool invariant $\text{inv}_{t'}(\mathbf{i}, \tau[t : (C', r)], \mathcal{T}[t' : (\text{skip}, r')], v_{t'}, \Delta')$ in case of $t' \neq t$.

We now show that $\text{inv}_t(\mathbf{i}, \tau[t : (C', r)], \mathcal{T}[t' : (\text{skip}, r')], v_t', \Delta')$ hold, both when $t = t'$ and $t \neq t'$. Let us first assume $t = t'$ ($r = r'$). Then from (19) we get that $A = \mathcal{L}(m, a, r)$ and $\Delta'(t) = \text{done}(\mathcal{L}(m, a, r))$ hold. When $t \neq t'$, no abstract transition is made in t , so $\Delta(t) = \Delta'(t)$ and $\mathcal{T}(t) = \mathcal{T}[t' : (\text{skip}, r')](t)$. Consequently, the following is true in both cases:

$$\begin{aligned} & ((\Delta'(t) = \text{todo}(\mathcal{L}(m, a, r)) \wedge \mathcal{T}[t' : (\text{skip}, r')](t) = (\mathcal{L}(m, a, r), r)) \vee \\ & (\Delta'(t) = \text{done}(\mathcal{L}(m, a, r)) \wedge \mathcal{T}[t' : (\text{skip}, r')](t) = (\text{skip}, r'))). \end{aligned} \quad (23)$$

Together with $\text{safe}_t(v_t', C', \llbracket \mathcal{Q}(t, \mathcal{L}(m, a, r)) \rrbracket_{\mathbf{i}})$, those observations imply $\text{inv}_t(\mathbf{i}, \tau[t : (C', r)], \mathcal{T}[t' : (\text{skip}, r')], v_t', \Delta')$.

Combining these observations with $(\sigma', \Sigma', \Delta') \in \lfloor v_t' * \otimes_{k \neq t} v_k \rfloor$ following from (20), we conclude that we have demonstrated existence of v_1', \dots, v_N' such that:

$$(\forall k. \text{inv}_k(\mathbf{i}, \tau[t : (C', r)], \mathcal{T}[t' : (\text{skip}, r')], v_k', \Delta')) \wedge (\sigma', \Sigma', \Delta') \in \lfloor \otimes_{k \in \text{ThreadID}} v_k' \rfloor,$$

which by the induction hypothesis $\phi(n-1)$ implies that $h \in \mathcal{H}_{n-1}[\ell, \tau[t : (C', r)], \sigma'] \subseteq \mathcal{H}[\mathcal{L}, \mathcal{T}[t : (\text{skip}, r')], \Sigma']$. Now that we demonstrated that $\mathcal{T}(t') = (\mathcal{L}(m', a', r'), r')$, $\Sigma' \in \llbracket \mathcal{L}(m', a', r') \rrbracket_t(\Sigma)$ and $h \in \mathcal{H}[\mathcal{L}, \mathcal{T}[t : (\text{skip}, r')], \Sigma']$ all hold, by Definition 15 we can conclude that $h \in \mathcal{H}[\mathcal{L}, \mathcal{T}, \Sigma]$. \square